Sensory Sensitivity Wayfinding

Milestone 3 Final Report

iMOVE CRC Project #3-048

iMOVE project number and title	3-048 Sensory Sensitivity Wayfinding
Milestone number and deliverable name	Milestone 3 Final Report (user research findings, MVP feature specifications, iterative development process, evaluation results, and post-pilot recommendations)
Partners	Autism Spectrum Australia (Aspect); Public Transport South Australia (SA Department for Infrastructure and Transport)
Report Completed by	Keith Urquhart, MapitX trading as BindiMaps
Date of Report	10 June 2026

This research is funded by iMOVE CRC and supported by the Cooperative Research Centres Programme, an Australian Government initiative.

Table of Contents

Executive Summary **4**
1. Introduction **5**
1.1 Project origin and method 5
1.2 Relationship to Milestone 2 5
1.3 What this report covers 6
2. Project Context & Objectives **6**
2.1 Sensory sensitivity in the transport context 6
2.2 Existing accessible-wayfinding substrate 7
2.3 Milestone 2 findings recap 8
2.4 Existing sensory-sensitive wayfinding resources 10
3. User Research & Insights **12**
3.1 Overview 12
3.2 Method 13
3.3 Thematic findings 14
3.4 Contradictory preferences as a finding 21
3.5 Scope and transferability 22
3.6 Outcome 23
4. MVP Feature Specifications **23**
4.1 Product framing 23
4.2 Publishing workflow 23
4.3 Feature catalogue 24
4.4 Highlighted features 36
4.5 What was deliberately not built 45
5. Design Iteration **47**
5.1 Method 47
5.2 Pre-Focus-Group 1 build 47
5.3 Iteration breadth: Focus Group 1 to Focus Group 2 48
5.4 Iteration depth: five worked cycles 48
5.5 How risks were managed in iteration 51
6. Pilot Deployment **51**
6.1 Pilot site and partnership 51
6.2 Deployment 52
6.3 Operator-supplied dynamic content 55
6.4 Partner experience 56
7. Evaluation Metrics & Results **57**
7.1 Methods and instruments 57
7.2 Engagement 57
7.3 Stress and confidence change 58
7.4 Satisfaction 59
7.5 Quality metrics 59
7.6 Limitations 61

8. Data-Dependency Findings & Pivots..... 62
8.1 Originally intended dynamic data sources..... 62
8.2 Feasibility findings per source..... 62
8.3 Convergence with the user research..... 64
8.4 What was delivered as dynamic content..... 64
8.5 Lessons for transferability..... 65
9. Framework & Design Guidelines..... 65
9.1 The Sensory-Sensitive Wayfinding Framework..... 66
9.2 Design guidelines..... 70
9.3 Open framework questions..... 76
10. Recommendations for Broader Deployment..... 76
10.1 The validated set..... 76
10.2 Next-stage research..... 78
10.3 Australian standards gap..... 78
10.4 Platform and partner deployment..... 79
10.5 Data-availability recommendations..... 81
10.6 Post-MVP feature recommendations..... 81
10.7 Residual risks and continuity..... 85
11. Conclusion..... 86
11.1 Findings..... 86
11.2 Outstanding questions..... 87
11.3 Acknowledgements..... 87
Appendices..... 88
Appendix A. Milestone 2 reuse note..... 88
Appendix B. Audit and quality summary..... 88
Appendix C. Glossary..... 90
Appendix D. References and prior-art list..... 93
Appendix E. Demonstration videos..... 98

Executive Summary

Sensory sensitivity is atypical reactivity to everyday sensory input (sound, light, smell, touch, movement, crowding) at a level that disrupts daily participation. Most prevalent in autism, it also occurs across ADHD, anxiety, PTSD and dementia, and in a significant minority of the general population.

The Sensory Sensitivity Wayfinding project explored how digital wayfinding tools can help people with sensory sensitivities at complex public transport hubs. Funded by iMOVE CRC, BindiMaps led the project and brought in Autism Spectrum Australia (Aspect), one of Australia's largest autism-specific service providers, as research partner. The pilot ran at Adelaide Railway Station with Public Transport South Australia (PTSA), the customer-facing public-transport authority under the South Australian Department for Infrastructure and Transport, as operator partner. The live pilot guide can be viewed at https://sensory-guide.web.app/venue/adelaide-railway-station.

Across four rounds of testing, the project pivoted from sensory map to planning guide. Participants treated the resource as something to read before a trip rather than consult during one. The original sensory map overwhelmed them. The redesigned guide leads with progressive disclosure, real photographs and multi-select sensory filters, with the map as a secondary reference. The redesign was received positively in the second focus group, with overwhelm language absent.

From a single published draft, the guide produces five outputs: an interactive web page, an embeddable partner version, a filterable PDF, a Word document and an audio narration. Live transit-crowding, IoT sound, maintenance and popular-times feeds were not available for Adelaide Metro. The admin tool replaces them by letting venue staff or auditors edit the guide directly, with the public copy refreshing against the same URL.

Accessibility and security reviews of the build, combining automated scanning with manual inspection, cleared the guide for continued deployment subject to two short remediations on the security side. Independent third-party audits are recommended before wider deployment (§10.4). The work produced a six-component framework with 22 design guidelines, each traced to focus-group evidence and to a feature built or flagged for future build.

The top recommendation: run a pre-and-post planning-anxiety and confidence survey alongside a wider deployment. Test the same approach at other venues and with other sensory-sensitive groups.

1. Introduction

1.1 Project origin and method

Sensory sensitivity is atypical reactivity to everyday sensory input (sound, light, smell, touch, movement, crowding) at a level that disrupts daily participation. Standardised measures detect atypical sensory-processing patterns in around 95% of autistic children compared with typically-developing peers (Tomchek and Dunn 2007), and the pattern recurs across ADHD, sensory processing disorder, anxiety, PTSD and dementia. Roughly one in five children in the general population also show at least one atypical sensory pattern (Dean et al. 2022).

This project began as research into what makes built environments hard for people with sensory sensitivities. Milestone 1 (2024–25) was problem-scoping: semi-structured interviews with people with sensory sensitivities and the practitioners who support them. That round surfaced noise and auditory load as the dominant sensory channel, alongside lighting, crowding and unfamiliarity. Predictability emerged as the underlying access need, drawn out through synthesis with Aspect's public guidance and the social-model literature.

BindiMaps' existing accessible-wayfinding platform offered one credible vehicle for addressing those problems. The platform was built first for users who are blind or have low vision and generalised outward, with three relevant foundations: a maintained indoor-mapping substrate, an audit-led content workflow, and an accessibility-first product team.

The method across all three milestones: research, propose, test, correct. The largest correction was the map-to-guide pivot. The original design was a sensory-mapping interface. Milestone 2 testing and Focus Group 1 showed that the map format increased anxiety rather than reduced it. The product was rebuilt as a planning-oriented Sensory Guide with maps as secondary reference.

1.2 Relationship to Milestone 2

Milestone 2 produced the Static Sensory Mapping Report: annotated site maps of Adelaide Railway Station, a mapping methodology, sensory-trigger categories with severity levels, and a validation focus group on the static product. Six findings carried into Milestone 3, set out in full in §2.3.

Milestone 3 built on those findings: an interactive product, lived-experience review retained across iterations, the map-versus-guide question tested against successive working prototypes at the same venue.

1.3 What this report covers

The report covers research, prototyping, pilot deployment and evaluation of the Sensory Guide at Adelaide Railway Station with PTSA as operator partner. Six lived-experience participants contributed across the two focus groups, alongside three auditors from Autism Spectrum Australia (Aspect) who tested the admin tool. Aspect ran the audit, recruited the focus-group participants, designed the question sets, facilitated the sessions and reviewed each iteration cycle. The Aspect Autism Friendly team facilitators hold both professional roles and lived-experience perspective.

The report does not attempt statistical generalisation across the broader sensory-sensitive population. The cohort is small and purposively recruited through Aspect, the venue is single, the deployment window is short. Where the evidence supports a finding, the report states it. It does not estimate the size of the effects it finds, nor test whether the same patterns would hold at other venues. Both are recommended in §10. The depth-over-breadth trade-off is deliberate. Two of the findings were also cross-checked against the larger Buro Happold survey of 280 respondents (§3), where the survey question lined up.

2. Project Context & Objectives

2.1 Sensory sensitivity in the transport context

Autistic people and people with sensory processing differences face significant barriers in public transport environments. Noise, lighting, visual complexity, crowding and spatial transitions each raise stress and erode the confidence needed for independent travel, and they are rarely experienced one channel at a time. Cumulative load, more than any single trigger, determines whether a trip is feasible.

At Adelaide Railway Station, the Milestone 3 pilot venue, the Aspect audit (Appendix A) records the kind of point-source noise that recurs across the south-side concourse: ‘There is a loud metal clunk sound at the start of the escalator that rises up to ground level.’ The same audit documents a deep rumble from trains overhead and bright, uncomfortable downlighting at the tunnel ends.

The peer-reviewed literature corroborates this pattern. Haas and colleagues, working with 14 young Autistic adults in Sydney, identified three interacting factors in transit stress: intolerance of uncertainty, the dominant role of anxiety, and sensory processing, particularly crowding and the tactile, auditory and visual stimuli that accompany it (Haas et al. 2020, p. 5). The same study recorded the anticipatory dimension that an in-venue audit misses:

> *‘It makes me really quite anxious that I will miss my stop since I do not know where the bus is going, and I don’t know which stop to look out for, even if I’ve looked it up*

> *online. I remember avoiding using buses for quite a few years as I did not feel comfortable finding the right stop, since bus routes are less obvious than trains, and you have to ask them to stop.'*
>
> *(Participant YA02, Haas et al. 2020, p. 16)*

Neil and colleagues (2016) quantified the relationship in Autistic children: intolerance of uncertainty and anxiety accounted for approximately half the variance in their sensory sensitivities, but only around a fifth in typically developing children. Autistic adults report the cumulative quality of the load in their own words. One participant in MacLennan, O'Brien and Tavassoli's 2021 study put it this way:

> *'[…] all the loud environmental noise (which may seem like nothing to most people—air conditioning humming, projector buzzing, lights buzzing, plates clinking in a restaurant kitchen), can drive me into a shutdown.'*
>
> *(Participant SE040, MacLennan et al. 2021)*

Three implications follow for transport access.

- **Predictability is an access requirement, not a preference.** Advance information about what a venue contains determines whether the visit is feasible at all. Advance information about whether the lift is working only routes around a known barrier on a trip already committed to. Aspect, in its public guidance for the arts sector, states that 'unpredictable environments increase anxiety and reduce participation. […] Predictability does not reduce enjoyment. For many Autistic people, it increases it' (Aspect, *Access in the arts: What makes a space autism-friendly*, accessed 19 May 2026).
- **Anticipatory stress, not only in situ.** Sensory-sensitive travellers plan around expected load before they leave home. Without usable pre-trip information, the trip may not be attempted at all.
- **Structural barrier, not individual.** The social model of disability locates the disabling effect in 'a contemporary social organisation which takes no or little account of people who have physical impairments and thus excludes them from participation in the mainstream of social activities' (UPIAS 1976). On that view, a station that publishes nothing usable for planning is itself the disabling element.

Within that frame, Milestone 1 research (2024–25) formulated four problem hypotheses for sensory-sensitive transport access: noise and auditory load (the anchor channel), lighting, crowding and unfamiliarity. The Milestone 3 build addresses all four together.

2.2 Existing accessible-wayfinding substrate

BindiMaps' existing accessible-wayfinding platform was one credible vehicle for addressing the Milestone 1 problems. The platform was built first for users who are blind or have low vision,

then generalised to users of all abilities. Building for the highest-need cohort first, rather than retrofitting accessibility onto a mainstream product later, left useful foundations in place for sensory-sensitive wayfinding.

That approach has theoretical and empirical backing in three overlapping bodies of work.

- **Universal design (Mace 1997)** treats accessibility as a property of the artefact rather than an accommodation layer added later.
- **Inclusive design's persona spectrum (Microsoft 2016)** reframes capability across permanent, temporary and situational modes: a feature that helps an Autistic traveller plan a noisy commute also helps a person experiencing a migraine or a jet-lagged traveller manage the same environment.
- **The curb-cut effect (Hamraie 2017)** is the canonical name for the pattern in which accessibility features designed for disabled users deliver disproportionate benefit to the non-disabled majority.

The Sensory Guide, the resulting MVP, is a new web application built for this project, with BindiMaps' existing accessible-mapping products integrated into it.

2.3 Milestone 2 findings recap

Milestone 2 (2025) was the static-mapping deliverable in the project's pre-prototype phase: it produced the Static Sensory Mapping Report, comprising annotated site maps, the mapping methodology, sensory trigger categories and user feedback from a validation focus group. Six findings carried into Milestone 3.

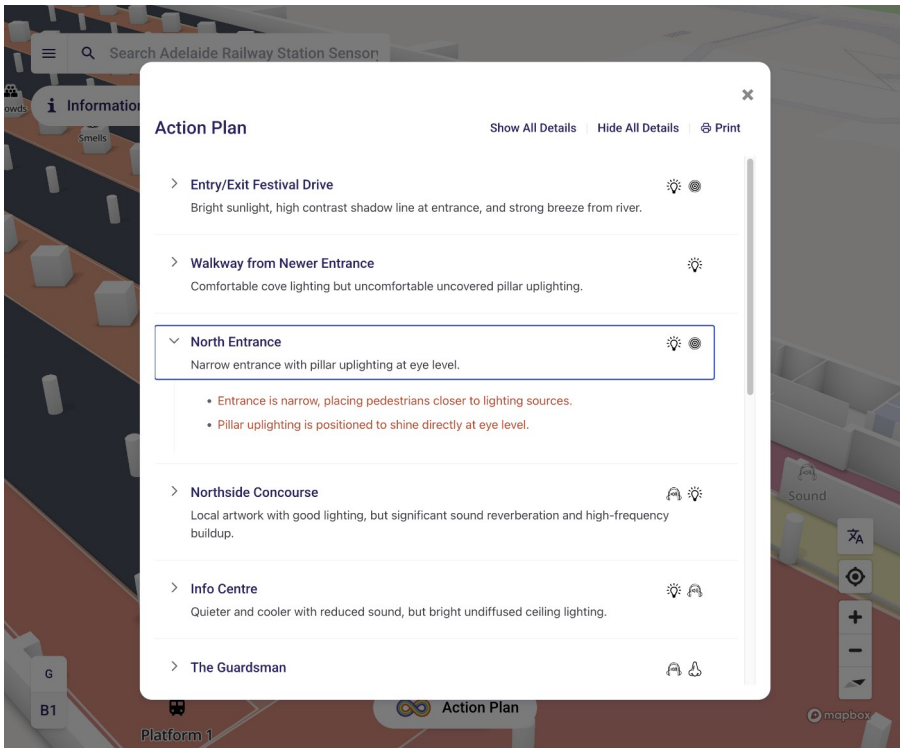

Milestone 2 action-plan prototype tested alongside the sensory map. Focus-group preference for the action-plan format gave the early signal for the Milestone 3 pivot.

- The categorisation framework (sensory-trigger categories and severity levels) covered Adelaide Railway Station's features and was clear to lived-experience reviewers.
- Aspect and the Milestone 2 focus group changed category definitions and severity wording before the framework was used to author content. Lived-experience reviewers shaped the framework upstream.
- Reviewers preferred action-oriented content (what to do, what to expect, where to go) over neutrally-descriptive content.
- Sensory challenges were experienced cumulatively rather than as isolated triggers. Areas where multiple inputs co-occurred were consistently identified as more challenging than any single input on its own.
- Static sensory information was valued for preparation, even where triggers could not be avoided. Advance awareness supported confidence and reduced uncertainty in planning.
- Reviewers' early signals that a map-led approach raised anxiety rather than reducing it foreshadowed the Milestone 3 pivot away from a map-centric interface.

These findings set the Milestone 3 scope: build the Sensory Guide, retain Aspect-led lived-experience review across iterations, and test the map-versus-guide question against successive working prototypes at Adelaide Railway Station.

2.4 Existing sensory-sensitive wayfinding resources

The existing corpus is dominated by static PDFs, distributed across museums, airports, shopping centres, local government, arts venues, parks and a small number of stadiums. Australian airports have effectively standardised the format: Hobart, Adelaide, Gold Coast, Cairns, Townsville, Melbourne and Canberra each publish a per-terminal sensory map PDF and a sensory or social story PDF, with variants of the Hidden Disabilities Sunflower programme and onsite quiet rooms appearing at several. Museums and galleries publish visual stories and sensory maps with similar consistency. The Stockland network publishes PDF guides of the same shape across shopping centres. Local-government libraries and community centres publish social-story PDFs.

The corpus breaks down by sector as follows.

Sector	Dominant format	Coverage	Representative examples
Airports	Per-terminal sensory map and social-story PDF	Standardised	Hobart, Adelaide, Gold Coast, Cairns, Townsville, Melbourne, Canberra
Museums and galleries	Visual story and sensory map (mostly PDF, some web)	Consistent	Australian Museum, NGV, WA Maritime Museum; Tate Modern and V&A (UK)
Shopping centres	Sensory guide and map (web plus PDF)	Consistent	Stockland network, DFO Brisbane
Arts and performing arts	Visual-story library; per-production guides	Emerging	Arts Centre Melbourne, MCG
Local government, libraries and civic	Social-story PDF; sensory map	Emerging	Monash Libraries, Georges River Council
National parks and outdoor	Multi-format (PDF, editable Word, Key Word Sign)	Emerging	Parks Victoria, Zoos Victoria

Sector	Dominant format	Coverage	Representative examples
Rail and metro	Visual stories	Emerging	Sydney Trains and Sydney Metro
Sport and events	Sensory map	Sparse	Australian Open
Education and research	Sensory map; sensory discovery space	Sparse	University of Wollongong
Healthcare	None	None verified	No Australian hospital met the verification bar (§10.2)

The full catalogue with sources is at Appendix D.

Three patterns from this corpus informed the Milestone 3 design.

- **Web-first and interactive formats are the exception.** The Australian Museum's visual story, Tate Modern's visual story, the MCA's web sensory map, Monash Libraries' web-hub model and DFO Brisbane's web-plus-PDF approach are notable precisely because most of the corpus is downloadable PDF.
- **Sub-resource granularity is emerging.** Per-exhibition resources (NGV; the National Museum of Australia's *Mr Squiggle* sensory map), per-event-type stories (the MCG's separate footy and cricket sets), per-production guides (the Arts Centre Melbourne visual-story library; Stratford Festival) and directional sets (Canberra Airport's departure/arrival × domestic/international story matrix) appear across the more recent material.
- **Three certification frameworks operate without cross-recognition.** Aspect's Autism Friendly Accreditation (audit-led, scoped to certified Quiet Rooms with twice-yearly re-audit), the Autism SA Autism Friendly Charter (training-led, with a 75% workforce-training threshold and a public business directory) and Dementia Australia's design principles (Fleming and Bennett, University of Wollongong) coexist in Australia. None recognises another.

Healthcare is the standout gap. No Australian hospital sensory map, social story or action plan met the verification bar during prior-art review. Hospital passports are referenced in policy material but no public copies could be located.

3. User Research & Insights

3.1 Overview

User research ran across four sessions from Milestone 1 to Milestone 3. The cohort was Autistic adults with lived experience of sensory sensitivities. Autism Spectrum Australia (Aspect) recruited the participants and organised the sessions. Aspect and BindiMaps jointly designed the question sets and ran each session. Problem-research started in Milestone 1, asking what makes built environments hard for people with sensory sensitivities, before any commitment to a particular product form.

The three later sessions were the Milestone 2 sensory-map validation focus group, a Milestone 3 focus group on the sensory-map prototype (Focus Group 1, 24 March 2026), and a Milestone 3 focus group on the redesigned Sensory Guide (Focus Group 2, 19 May 2026). Aspect reviewed each iteration cycle between sessions. The findings drove every meaningful product decision and underwrite the validated design principles in §9.

User-research arc, Milestone 1 to Milestone 3

Four research sessions of escalating fidelity. Aspect's Autism Friendly team recruited participants, co-designed the question sets and ran every session.

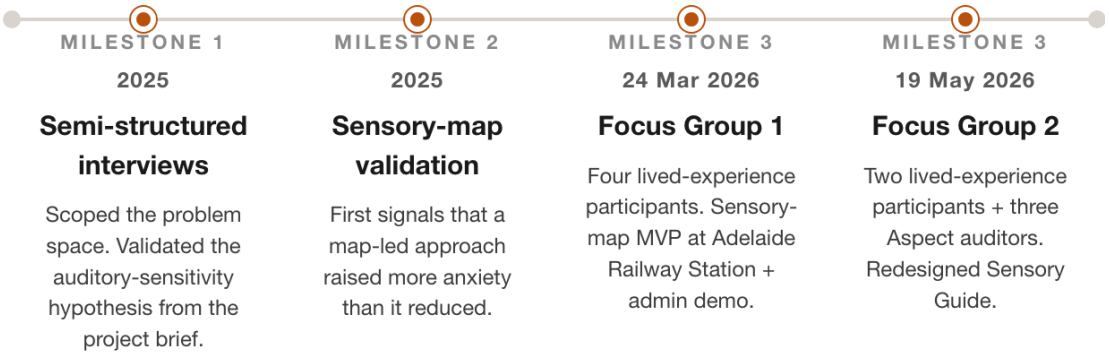

ASPECT REVIEW BETWEEN ITERATIONS Aspect supplied lived-experience review on each iteration without waiting for a focus group, allowing iteration on smaller changes between sessions.

3.2 Method

Research arc

The four sessions deliberately escalated in fidelity.

- **MS1 interviews (2024–25).** A series of semi-structured interviews with people with sensory sensitivities and the practitioners who support them. Purpose: scope the problem space and test the noise-and-auditory hypothesis carried into Milestone 1.
- **MS2 sensory-map validation (2025).** A focus group reviewing the initial sensory-map concept. Surfaced the first signals that a map-led approach raised more anxiety than it reduced for the target users.
- **MS3 Focus Group 1 (24 March 2026).** A three-hour remote session with Aspect's Autism Friendly team and four lived-experience participants. Three-part structure: an agree-disagree walkthrough of earlier findings, hands-on exploration of the sensory-map MVP at Adelaide Railway Station, and a demo of the admin ingestion flow.
- **MS3 Focus Group 2 (19 May 2026).** A two-hour session split in two halves. Part one tested the redesigned Sensory Guide with two Autistic participants who are not Aspect staff. Part two tested the redesigned admin tool with three Aspect auditors as prospective admin users.

Aspect's role

Aspect's Autism Friendly team consults with airports, shopping centres, stadiums, councils and transport operators across Australia, including Sydney Trains. Its staff bring both professional and lived-experience perspective. For this project Aspect:

- conducted the Adelaide Railway Station sensory audit (Appendix A)
- recruited participants for both focus groups
- co-designed the question sets
- ran the sessions
- reviewed each iteration cycle in between

Cohort

Across the Milestone 3 sessions, six unique lived-experience participants contributed across the two focus groups, with two Aspect facilitators in a separately-noted dual role.

- **Focus Group 1**: four lived-experience participants (FG1-P1 to FG1-P4). Three were Aspect Autism Friendly team staff participating in lived-experience capacity; one was external to Aspect.

- **Focus Group 2**: two external lived-experience participants (FG2-P1, FG2-P2) tested the redesigned Sensory Guide.
- **FG2 admin testing**: three Aspect auditors tested the admin tool in a separate Focus Group 2 session (FG2-A1, FG2-A2, FG2-A3). Two of the three had attended Focus Group 1 in auditor capacity.
- **Aspect facilitators**: two Aspect Autism Friendly team staff co-designed the question sets and ran both sessions. They also contributed lived-experience commentary in session, coded FG1-P5, FG1-P6, FG2-P3, FG2-P4 when quoted in that capacity. They are noted separately from the six lived-experience participants above.

Participants are referred to by code throughout.

3.3 Thematic findings

Seven themes emerged across the two focus groups. They are not ranked. Each shaped a product decision. Quotes are lifted from the session transcripts verbatim, lightly cleaned of speech disfluencies, with elisions marked `[…]` and inserted words in `[brackets]`.

Choice and control

Avoidance is the obvious response to a difficult sensory environment, but preparation can work just as well. The same stimulus that overwhelms in one context can be welcomed in another, depending on whether the person chose to be there and knew what was coming.

In a Milestone 1 interview, MS1-P1 described shopping centres as difficult because the sound is layered music, voices and noise without warning. At a live concert she will stand in the front row of a mosh pit and block out the noise for the set, because she chose to be there, she knows what is coming, and she is there for the performer. Without that anticipation and choice, she said, the same volume would have her on the floor.

The Milestone 3 focus groups carried the same pattern:

> *‘Some days I am prepared to go into the bathroom and I’ll be all right. I’ll wear my headphones, and [as for] the noise of the hand dryers, I’ve been expecting it, so I can manage that today. […] But other times that same person, I don’t want to even go near the hand dryers.’ (FG1-P3)*

Same person, same stimulus, opposite outcomes. The product cannot remove the sensory environment, but it can return anticipation and choice to the visitor before they encounter it. Framed this way, personalisation features (filters, collapsible sections, opt-in toggles, multi-format export) are access requirements. Each one returns to the visitor a decision the product would otherwise have made on their behalf.

An FG1 participant put end-user control concretely:

> *‘I love that you could filter by the senses. That was really helpful. So I just put the key senses that I’m impacted by and went through by those.’ (FG1-P1)*

An FG2 participant said the same about the redesigned category filters: ‘I like the categories. I think that’s quite a good representation of it.’ (FG2-P1). FG1-P5 named the principle the project adopted as the design’s organising label: ‘it’s about choice and control […] being able to control your experience and choose what information you want.’

The same principle holds on the admin side, in the opposite direction. The Focus Group 1 admin tool used AI to rewrite an Aspect auditor’s PDF into structured guide entries, and participants strongly rejected the arrangement. Aspect staff are the domain experts on their own assessments, and AI silently restructuring those words removes their authorship of what gets published. Visitors must retain control of what they see, and admins must retain control of what they publish. See *AI as helper, not author* below for the detail and quotes.

Planning over in-moment use

Every FG1 participant treated the product as a pre-trip planning aid rather than a live-navigation tool, unprompted within minutes of the session opening.

> *‘I love the action plan because I spend a huge amount of time planning my journey, with a large emphasis on avoiding overwhelming areas and finding quiet spots. This plan helped me to do this.’ (FG1-P1, chat)*

FG1-P1 on the mechanism:

> *‘It reduces my anxiety to be able to picture and anticipate what’s going to happen. So if I can look at photographs and imagine myself there […] it takes away a lot of the anticipatory anxiety.’ (FG1-P1)*

FG1-P6, in lived-experience capacity, framed her own preparation as an energy budget:

> *‘there’s only a limited amount of energy I have, so I’m planning, planning around conserving that.’ (FG1-P6)*

An FG2 participant generalised the pattern outside the rail context:

> *‘When I was a kid and I would catch a bus for the first time somewhere, I would have to know every bus route, every timetable. […] it reduces my anxiety and makes me feel better about being out in public because it reduces the layer of unpredictability.’ (FG2-P1)*

The same planning need shows up in independent prior art at much larger scale. Buro Happold’s 2024 survey of 280 people asked how sensory maps help. One respondent answered:

> *'Sensory maps allow me to plan better, to understand the environment I would be going into, and avoid areas that would cause overload.'*
>
> *(Survey participant, 2024. Buro Happold, Sensory Mapping (2026), p. 5)*

Across both focus groups no participant described wanting in-the-moment help over advance-of-trip planning.

Within planning, the focus group raised a design choice. The intuitive move is to engineer the problem away: an app that automatically routes the visitor around busy or stimulating areas, treating sensory categories the way a navigation app treats traffic. This is a recurring proposal in venue-accessibility design and a natural extension of map-based wayfinding.

The alternative is to simply prepare the visitor for what they will encounter, rather than over-engineer a solution. It is more pragmatic, cheaper to build, and provides almost the same value to the user, arguably more.

FG1-P5 spelled out the practical dimensions: forewarning alongside in-moment coping, places to recover and broader environmental predictability. She noted that forewarning helps but does not eliminate the impact: *'knowing about things beforehand reduces the distress and overwhelm, it doesn't eliminate it'* (FG1-P5, transcript 44:59).

The guide cannot change the venue. It can support in-moment coping only indirectly, by flagging quieter areas. What it can do reliably is forewarn the visitor about what to expect.

FG1-P1 on the corresponding visitor experience:

> *'Interesting thought. I would say ideally I don't want the big sensation. BUT if it's unavoidable, knowing it will happen will help to reduce the impact. I think it's important to give [people] the information so that they can make a decision to go or not to go, depending on their capacity and sensory needs that day.' (FG1-P1, chat)*

Two implications follow. First, forewarning lets the visitor decide for themselves whether to go and how to prepare, which the participants treated as the preferred form of help. System-driven avoidance promises something a venue guide cannot reliably deliver: the guide cannot anticipate unannounced events, crowd surges or last-minute changes, and any silent inaccuracy is invisible to the visitor relying on it.

Second, cost and effort. Forewarning is a publishing problem, addressable by giving a venue a structured place to record what is true of their space. Auto-routing around sensory categories is a substantially larger build for less reliable mitigation.

The pivot from a sensory map to a Sensory Guide followed from this finding.

Photographs over schematic maps

Both focus groups preferred real photographs of the venue over the map. Photos let participants picture themselves in the space. The map was the more abstract representation, and took more cognitive work to interpret.

> *'I really enjoyed having the photo so I could picture myself there.' (FG1-P1)*

360-degree walkthroughs emerged unanimously as a complement to still photos. A still image fixes one vantage point, whereas a walkthrough lets a visitor turn to orient themselves and plan a route through the space before arriving, which a single photograph cannot do. Both sessions also converged on the same shortfall in the MVP: the public guide carried roughly one photo per area, and the cohort signalled the need for significantly more.

AI as helper, not author

The admin tool tested in Focus Group 1 used an AI model to read an Aspect audit PDF and rewrite the content into structured guide entries. The human auditor's role was to review the AI's output at the end. Participants rejected this arrangement and named a different relationship between human and AI: AI is welcomed as a helper that suggests and flags, but rejected as an author that silently rewrites.

FG1-P3 was direct about the discomfort:

> *'I know some people in our team use AI for different reasons. Personally, I try to avoid it wherever I can. It makes me feel very uneasy for different ethical reasons.' (FG1-P3)*

FG1-P1 named the helper role the group wanted:

> *'I think if it we were cautious of it not changing language, but even just identifying, hey, we've noticed an error, would you like it to change it to this might be really good.' (FG1-P1)*

FG1-P5 flagged non-determinism as a trust issue:

> *'as it's pulled out different information, it looks different than the first one that you've developed that we tested […] every time you upload the same document, it would work differently' (FG1-P5)*

FG1-P5 then named the helper role: 'I think sometimes it can be helpful to check our work rather than changing our work and us checking it.' And again: 'while tech is helpful, it doesn't accurately describe the experience of autistic people […] the information needs to be developed based on an assessment of the venue by people, by autistic people.'

The admin tool was rebuilt around this principle between FG1 and FG2. FG2 tested the rebuilt tool only, so this is the reception of the redesign rather than a controlled re-test. AI-anxiety language did not reappear in discussion of the new admin flow, and one of the returning auditors welcomed the change directly in chat:

> *'I really like that you can bypass the AI bits and upload your own words. This is great.' (FG2-A2, chat)*

Cumulative sensory load

The sensory-map prototype tested as a source of overwhelm. FG1 participants were direct about this.

> *'I did struggle to understand all of the information, and I think in part it was because there was competing things going on. So my pattern recognition was kind of jumping between things and not quite sure on like the hierarchy of what to start with first and follow to.' (FG1-P1)*

FG1-P3 described it as new vocabulary on top of existing sensory load: 'trying to work that all out at the same time with maybe not enough context of what different colours indicate or different intensities indicate […] makes it really complex.' (FG1-P3)

FG1-P5 on the same problem: 'a sensory map currently could be extra confusing because there's a lot of visual information. Some of it is a new language to learn.' (FG1-P5)

FG1-P1 again: 'I thought the information was great, it was just kind of too much at once and then I was overwhelmed and then I couldn't really like piece it through.' (FG1-P1)

The map was contributing to the problem it had been built to solve. This triggered the map-to-guide pivot. The redesigned guide tested in FG2 carried three changes:

- **progressive disclosure**: sections start collapsed, and visitors tap to expand only the parts they want
- per-category filters
- per-area 'Show on map' buttons

FG2 did not re-test the original sensory map, so this is the reception of the redesign rather than a controlled before-and-after. Overwhelm language did not reappear in discussion of the redesigned guide.

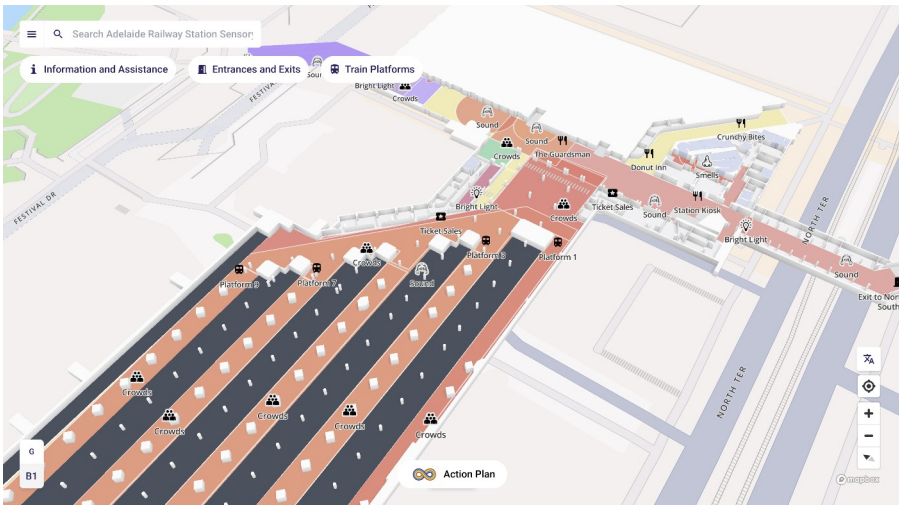

Sensory-map prototype that FG1 participants described as overwhelming

Familiarity of UI conventions

FG1 participants already knew Google Maps as a universal mental model, and bespoke interaction patterns paid a usability tax. The interface was rebuilt around the familiar convention.

FG1-P1: ‘Google Maps is just the software I tend to use, so I naturally kind of went to use the software as if it was Google Maps and I did get lost in it and get confused at different points.’ (FG1-P1)

FG1-P6 paraphrased the room’s consensus: ‘designing as close to the familiar as possible would be helpful.’ (FG1-P6)

FG1-P2 flagged a rotation-unrecoverable bug on the map: rotating worked, unrotating did not. FG1-P5 hit the same problem independently.

FG1-P5 also raised a terminology concern that extends beyond UI convention into language: ‘the naming convention for accessible routes […] I just think maybe in the future accessible routes won’t be the most appropriate term because accessible means different things for different disabilities. So I would probably say like mobility access route.’ (FG1-P5)

The rebuild delivered:

- collapsible sections
- an inline ‘Show on map’ button per area rather than a separate map mode
- standard scroll-and-tap mobile interactions

That rebuild removed the separate map mode the rotation bug lived in, so its absence in FG2 partly reflects a changed artefact rather than a fixed interaction. The rotation-recovery and map-mode confusion did not reappear in FG2, and the per-area ‘Show on map’ button was praised without prompting (‘Massively, yes’ (FG2-P1); ‘I think that’s a great addition’ (FG2-P4)).

Multi-format delivery

Multi-format publishing across web, PDF, Word document and audio guide became essential to the project. Two factors shaped this.

Government and enterprise procurement constraints prevent some venues from accepting a third-party-hosted page on their site, with PTSA as the trigger case. Individual participants also varied in their preferred consumption format.

FG2-P3 in the main session: ‘often, yeah, we do need to provide information in different ways or gather information in different ways. And so being able to have that flexibility and options is really valuable.’ (FG2-P3)

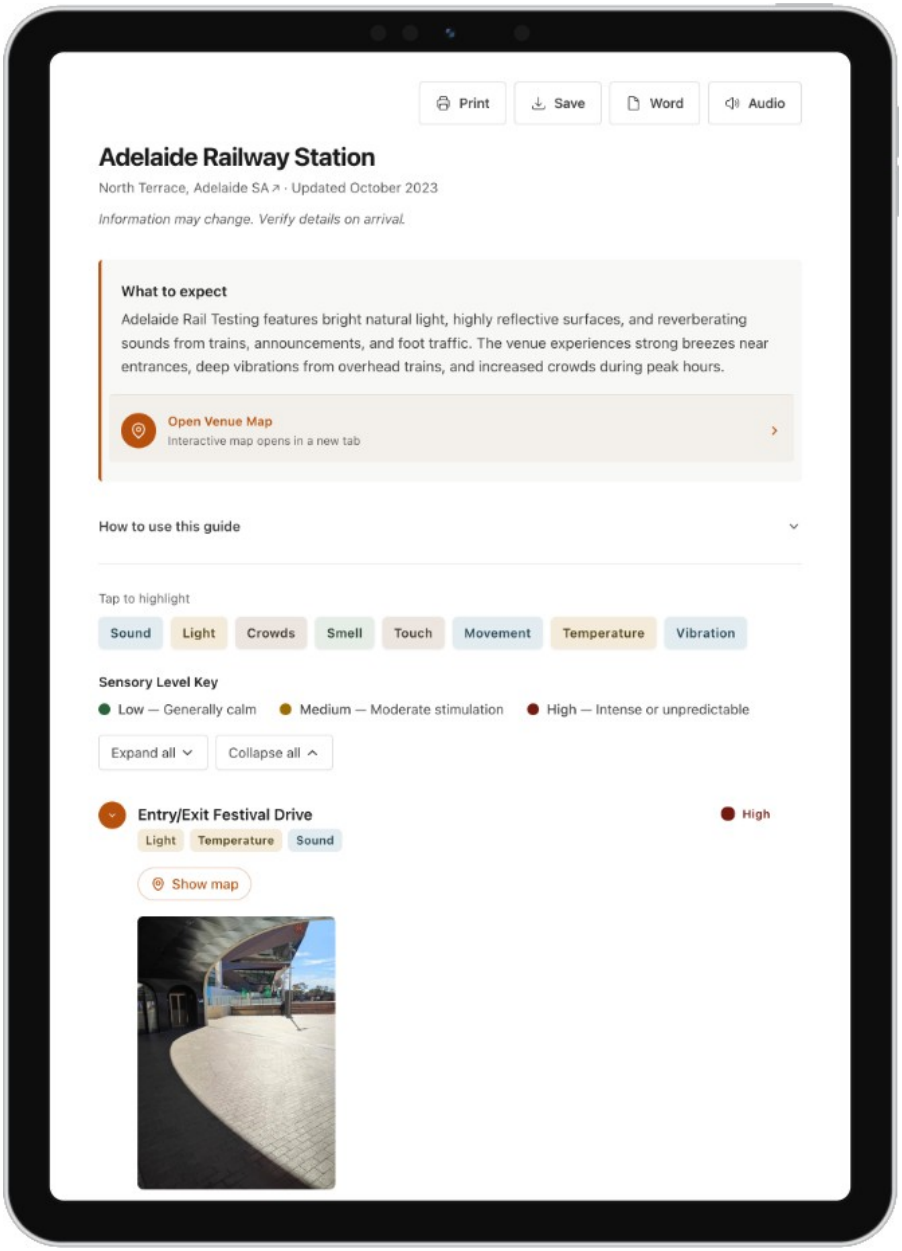

FG2 Sensory Guide with category filters, show-on-map buttons and collapsible how-to-use panel

3.4 Contradictory preferences as a finding

Across both focus groups, participants disagreed with one another on substantive design questions, and the disagreement is itself the finding. In accessibility work, opposite needs cancel one another out if a design picks a single position.

On information density. FG2-P1 favoured maximum detail: ‘I’m definitely the more information better version.’ (FG2-P1). FG2-P2 in chat, on the same overview: ‘I don’t like it, it mostly makes this go this isn’t for me. Also a lot of it is obvious I already know train station is going to be awful.’ (FG2-P2, chat). Two Autistic participants, same prototype, opposite conclusions about

how much content the page should carry. FG2-P1 then proposed the resolution: '[FG2-P2]'s idea of collapsing these sections might be a good balance between the two.' (FG2-P1)

On the primary entry point. FG1 was split between map-anchored and list-anchored entry points. Both camps signalled the other format should remain present. The disagreement was about which one greets the user first.

On planning versus in-the-moment use. FG2-P1 described heavy advance planning. FG2-P2 described arriving and looking it up on the spot: 'more likely when I arrive, unless the venue has told me this exists.' (FG2-P2). Both used the product, but each preferred a different default behaviour.

On the 'avoid versus prepare' axis. Even a single person's preferences shift across days. The hand-dryer description from FG1-P3 under 'Choice and control' shows the same stimulus moving between manageable and intolerable depending on whether she had prepared for it. FG1-P1 in chat: 'Ideally I don't want the big sensation. BUT if it's unavoidable, knowing it will happen will help to reduce the impact.' (FG1-P1, chat)

The design response to contradictory cohort preferences is personalisation: filters, opt-in toggles, collapsible sections, multi-format output and progressive disclosure to serve a cohort whose needs differ in this way. The alternative, a single default optimised for one end of the cohort, would underserve the other end.

3.5 Scope and transferability

This is qualitative work with a small purposive sample. It offers transferability to similar contexts rather than statistical generalisability.

Three scope boundaries shape how the findings transfer.

Aspect-led recruitment delivered deep Autistic lived-experience coverage at the cost of under-representing Deafblind, dementia-affected, PTSD, generalised anxiety, hearing impairment, non-autism sensory processing differences and older adults. Extending coverage to those constituencies is recommended in §10.

Both Milestone 3 sessions were anchored on a working prototype at Adelaide Railway Station, so venue-specific feedback such as station acoustics and platform navigation may not transfer cleanly to airports, hospitals or civic spaces.

Aspect's role across this research spans recruitment, question-set co-design, session facilitation, lived-experience review across iterations and admin-tool testing in Focus Group 2. Aspect also raised the possibility of a formalised audit-and-publish service at the end of Focus Group 2 (§10.4). Partner-led recruitment and facilitation is standard in accessibility research and increases depth of input, but it does concentrate methodological dependency on a single partner organisation, particularly in interpretation of the admin-tool reception. The report names Aspect's involvement transparently in each section that draws on Aspect-led input.

3.6 Outcome

Each principle traces to the focus-group transcripts and to a specific build change made before the second focus group. They are validated for this cohort at Adelaide Railway Station; testing their transfer to other groups and venues is a §10 recommendation. The principles are set out in full in §9.

4. MVP Feature Specifications

4.1 Product framing

The product was originally scoped as a *Sensory Awareness Mode*, a map-led overlay onto BindiMaps' existing accessible-wayfinding platform. What was built and tested is called the **Sensory Guide**.

Milestone 2 testing and Focus Group 1 showed map-centric framing increased anxiety for the target users, and planning-oriented framing tested significantly better. The evidence and the iteration are set out in §3 and §5.

The MVP sits on top of BindiMaps' existing platform, which provides venue map data, authentication, and ongoing distribution after the milestone closes.

4.2 Publishing workflow

A published guide passes through four stages.

- **Capture.** An auditor walks the venue with a phone or tablet, working area by area through a guided wizard. Photos come from the device camera, long-form descriptions from voice or typed text, and sensory details from category-and-level pickers. Capture is mobile-first because the work is.
- **Draft.** Multiple editors work on the same draft in real time, with each other's changes visible as they happen. Suggested AI edits surface for accept-or-dismiss review.
- **Publish.** One click promotes the draft to a versioned public copy and exposes a stable public URL. The same action regenerates the downloadable artefacts. Renaming a venue after publish flags it as unpublished until re-published, so the live copy cannot drift from the venue's current name.
- **Read.** The public guide renders at the public URL. The same content is available for inclusion on a partner site, as a filterable PDF, as a Word document, and as an audio narration. Visitors choose the format. Feedback is captured in-page via a short thumbs-up / thumbs-down with optional comment.

Admin venue dashboard for Adelaide Railway Station. Edit, Publish, Settings and Feedback are the entry points to the publishing flow above.

4.3 Feature catalogue

4.3.1 Public-side features

#	Feature	What it does	User need addressed
P1	Per-venue published guide	A public landing page per venue rendering the full Sensory Guide as one scrollable page.	Pre-trip planning (§3, *Planning over in-moment use*).
P2	Venue header with last-updated	Shows venue name, address, contact links and the date the guide was generated.	FG1-P5: ‘people knowing when it was updated is really helpful’.
P3	About-this-guide overview	A 1–2 sentence summary at the top of every guide describing what to expect sensorially.	FG2 information-hierarchy finding (§3).
P4	How-to-use onboarding panel	Collapsible help block explaining filters, sections, severity, busy times, map and key facilities; collapsed after first dismiss.	FG2-P1 on initial-load density; FG2-P2 on repeat-visit friction.
P5	Sensory category filter bar	Tap coloured category badges to highlight matching sections; multi-	§3 *Choice and control* (FG1-P1 on filtering by sense).

#	Feature	What it does	User need addressed
		select; selection persists across sessions.	
P6	Sensory key / legend	Compact key explaining the severity colour-coding; supports per-venue label overrides.	§3 *Cumulative sensory load*; colour-only coding was unreadable.
P7	Expand-all / collapse-all	Two buttons control every area section simultaneously.	FG2-P2 on scroll fatigue when everything is open.
P8	Journey-ordered sections	Areas sort by author-defined order so the guide follows the physical path through the venue.	Matches Aspect audit walkthrough format.
P9	Per-section progressive disclosure	Each area starts collapsed; tap to expand into preview, photos, sensory details and busy times.	§3 *Cumulative sensory load*.
P10	Area preview line	A short truncated preview under each collapsed area heading.	FG2 hierarchy finding; visitors skim before committing.
P11	Overall area severity badge	Each area shows an aggregate sensory level derived from its details.	FG2-P2: ‘If they have a rating they used to decide I want to know it’.
P12	Category badges per	Uppercase, text-only	§3 finding that icons

#	Feature	What it does	User need addressed
	area	colour-coded badges per area header.	did not reliably decode without text.
P13	Sensory detail entries	Per-detail description with category label, level label and a short prose description.	§3 *Multi-format delivery* / objective-language finding.
P14	Walk-time between areas	An optional minute estimate inserted between adjacent areas.	FG1-P6 on travel time between zones.
P15	Busy-times block per area	Optional days, times and description shown inside the expanded area.	FG1-P1; FG2-P2 on peak-times placement.
P16	Per-section photo carousel	Each area carries a carousel of photos with descriptive captions.	§3 *Photographs over schematic maps*.
P17	Image lightbox	One-tap full-screen photo viewer with adjacent-photo navigation.	FG1-P1 wanted zoomed-out context; the lightbox returns control to the reader.
P18	Show-on-map per area	An embedded mini-map for any area that has been pinned on the venue map.	FG2-P1: ‘Massively, yes’; FG2-P4: ‘a great addition’.
P19	Open-venue-map link	A floating link opens the interactive venue map in a new tab.	§3 *Photographs over schematic maps* contradiction; maps present but secondary.

#	Feature	What it does	User need addressed
P20	Per-area embedded media	Optional embedded media (BindiWeb venue page, YouTube clip) inside an area.	FG2-P3 on captioned media.
P21	Key Facilities block	Bottom-of-guide list of exits, accessible bathrooms and quiet zones.	FG2-P1 on key-facility hierarchy.
P22	Thumbs-up / -down with comment	After viewing, visitors can submit a short feedback note.	Post-visit reflection loop.
P23	Print PDF with QR codes	One-click PDF generation; embeds QR codes for the web URL and per-area map links in place of iframes.	PTSA content-governance constraint (§6); FG2 format-diversity finding.
P24	Filtered PDF	Same PDF generator with an option to download only sections matching active filters.	§3 *Choice and control* extends to print outputs.
P25	Word document export	One-click Word document download of the guide content.	PTSA request from its Customer Strategy and Inclusion Lead (§6); Word required to fit their internal template.

#	Feature	What it does	User need addressed
P26	Audio guide download	Generates an MP3 narration of the guide using a synthesised voice.	FG2 format-diversity finding; equal weight to other export formats.
P27	Iframe-embeddable	Each guide is embeddable as an iframe on a venue operator's site.	Multi-format delivery for partners that can host third-party embeds.
P28	Responsive layout	Single layout adapts from phone to desktop; touch targets meet 44 × 44 px guidance.	Different users prefer different devices across pre-trip planning and in-venue use.
P29	Keyboard and screen-reader support	Skip-to-content link, screen-reader cues, visible keyboard focus indicators, reduced-motion preference support and high-contrast variants.	Accessibility baseline.
P30	Powered-by footer	Footer link out to BindiMaps that opens in a new tab so it works inside iframes.	Partner attribution and discovery.
P31	Per-event analytics	Tracks views, expansions, filter toggles, image views, map opens, external links, downloads and feedback submissions.	Engagement and feature-trigger evidence collection.

#	Feature	What it does	User need addressed
P32	Home / landing page	Public home page introduces the product and links to BindiMaps.	Direct-visit entry point distinct from iframe embeds on partner sites.

4.3.2 Admin-side features

#	Feature	What it does	User need addressed
A1	Email and Google sign-in	Authentication via password, with forgot-password reset, or Google account.	Secure access for Aspect and venue staff.
A2	Venue dashboard	Lists every venue the signed-in user can edit, with a create-venue entry point.	Multi-venue scope.
A3	Create-venue wizard	Three-step flow (name, editors, start option) routes to capture or PDF import.	Multi-admin onboarding.
A4	PDF import (optional)	Upload an Aspect-style audit PDF; the system extracts text, parses areas and details, and writes a draft ready to edit.	Aspect deliverables come as PDFs; retained as fallback (§5).
A5	PDF import options	Toggleable rules at import: simplify language, force Australian English, force objective	§3 *AI as helper, not author*; auditors choose how much the import touches.

#	Feature	What it does	User need addressed
		language.	
A6	Import progress and rate limit	Real-time progress states with retry visibility and a daily import limit (50 per user).	Trust through visibility; cost control.
A7	Image extraction from PDF	Detects images embedded in the uploaded PDF and attaches them to matching areas.	Reduces re-keying for Aspect PDFs that already carry photos.
A8	Mobile-entry home	A phone-sized landing per venue: area list, add-area input, last-edited timestamp, presence indicators.	§3 *AI as helper, not author*; direct on-site capture replaces PDF re-typing.
A9	Floating-action quick-capture	Two phone-corner buttons: camera (new area, jump to photo step) and microphone (new area, jump to summary with voice on).	FG1-P5 on capture-as-you-go.
A10	Area capture wizard	Four-step linear flow per area: photos, summary, sensory details, review, with a tappable progress bar.	FG2 admin feedback that next-buttons were initially hidden.
A11	Voice dictation on long fields	A microphone affordance on every long-form text field;	FG1-P1 on dictation; FG2-P4

#	Feature	What it does	User need addressed
		cloud transcription writes back into the field.	endorsement.
A12	Live recording status	While recording: normalised input-level indicator, on-screen status, stop button.	On-site capture in busy venues.
A13	On-device photo capture	Tap to open the device camera or photo library; uploads attach to the active area.	§3 *Photographs over schematic maps*.
A14	Photo reorder and caption	Photos can be dragged into order, captioned and deleted with confirm.	FG1-P6 surfaced the photo-deletion gap; FG2 reordering endorsement.
A15	Cross-area photo manager	Desktop tab to drag photos between areas, with an Unassigned holding bay.	FG2-A1 on photo-source flexibility.
A16	Per-area severity per category	Within sensory details: pick category, level (Low / Medium / High) and description; multiple per area.	Sensitivity-category model; FG2 ordering revisions.
A17	Auto-save draft	Every change writes back to the venue's draft; presence and last-updated-by are recorded.	Multi-admin collaboration.

#	Feature	What it does	User need addressed
A18	Real-time multi-admin presence	Heartbeat presence per area; editor avatars next to each area in the list.	§3 collaborative-authoring theme.
A19	Content-update toast	When another editor saves a change, a non-modal toast announces who changed what.	Same as A18, shown non-modally.
A20	AI suggestions panel	On-demand call returns 3 to 15 inline suggestions (simplify, subjective, missing, inconsistency) each with explanation, accept and dismiss.	§3 *AI as helper, not author*.
A21	Per-suggestion before / after	Suggestion cards show the original phrase and the proposed replacement with a strike-through diff.	§3 trust through transparency.
A22	Post-import suggestion list	The PDF-import path produces a flat list of document-level suggestions (missing venue contact, missing or incomplete address, missing peak-hours information, underrepresented categories) above the draft preview.	Content completeness; FG2-P3 on publish-readiness warnings.

#	Feature	What it does	User need addressed
A23	Generate / regenerate overview and facilities	One-tap regenerate rewrites the venue's about-summary and extracts exits, bathrooms and quiet zones from the current draft.	Auto-extracted facility lists; FG2-P3 flagged in-place editing as a gap (§10).
A24	Daily AI usage display	Shows remaining quota of the 50-per-user-per-day limit shared across PDF transform, suggestions, overview and audio.	Cost and capacity awareness.
A25	Venue map URL	Field to attach a BindiWeb or other embeddable venue-map URL; powers the open-venue-map link and pin-on-map.	Maps as secondary reference.
A26	Pin areas on the map	Embedded BindiWeb map with a per-area pin editor; each area can be tied to one or more map locations.	FG2-P1 on map-to-section linkage.
A27	Per-area embedded-media URLs	Attach BindiWeb deep-links and YouTube videos to specific areas.	Multi-provider map support; FG2-P1 on menu links.
A28	Embedded-media URL validation	Pasted URLs are validated for safe embedding; broken or disallowed URLs surface as errors.	FG2-P3 saw a broken YouTube clip.

#	Feature	What it does	User need addressed
A29	Publish-readiness blocks	Publish is disabled until every selected sensory category has both a level and a description; the reason is shown inline.	Prevents silent data loss at publish time.
A30	Live draft preview	Full preview identical to the public layout, rendered from the live draft.	FG1-P3 on confirming venue names before publish.
A31	One-click publish	Promotes the draft to a versioned public copy; the public URL becomes live.	Publish step in the workflow.
A32	Publish-success summary	Post-publish view with public URL, share controls and links to outputs (embed code, PDF, Word, audio).	FG2 format-diversity finding.
A33	Version history	Lists every published snapshot; prior versions can be previewed, made live (revert) and deleted.	Content governance after venue changes.
A34	Rename flags unpublished	Changing the venue's display name after publish flags it as unpublished until re-published.	Content-accuracy discipline for partners.
A35	Outputs tab	Single hub listing the live URL, iframe	Multi-format delivery.

#	Feature	What it does	User need addressed
		embed code, audio file, PDF and Word document.	
A36	Embed code copy	Generates a ready-to-paste iframe snippet for a venue operator's website.	PTSA iframe-on-government-site conversation (§6).
A37	Audio caching by version	Each published version's MP3 is cached so repeat downloads do not re-spend AI quota.	Cost discipline.
A38	Per-venue editors list	A per-venue editors list (cap of five); any editor can add or remove others.	Collaborative-authoring theme.
A39	Editor invite link and script	Adding an editor generates a sign-in / reset link and pre-fills an invite-message text the admin can copy.	FG2-A1 asked for email automation; this is the MVP step.
A40	Per-venue feedback view	Aggregates public thumbs feedback and verbatim comments, filterable by 1, 2 or 4 weeks or all time.	FG2 admin feedback-discoverability gap.
A41	Admin help and onboarding	In-app help dialog covers dashboard, venue hub, capture and super-admin screens.	FG2-P3 first-use observations.

#	Feature	What it does	User need addressed
A42	Super-admin allow-list	Internal screen managing the email allow-list that controls admin account creation.	MVP trust boundary; every admin known to the team.
A43	Global analytics view	Super-admin aggregate dashboard of usage across all venues.	Internal operations and engagement evidence.
A44	Delete venue	Last-editor-only destructive flow that removes the venue.	Lifecycle and test-venue cleanup.
A45	Responsive admin	The same admin works on phone (primary for capture) and desktop (primary for review and publish).	Matches the dual capture / cleanup workflow.

4.4 Highlighted features

The features expanded in this section warrant rationale beyond what the catalogue row carries: a central research finding, a design pattern transferable beyond this project, or a residual tension flagged for §10.

4.4.1 Public-side: sensory category filters

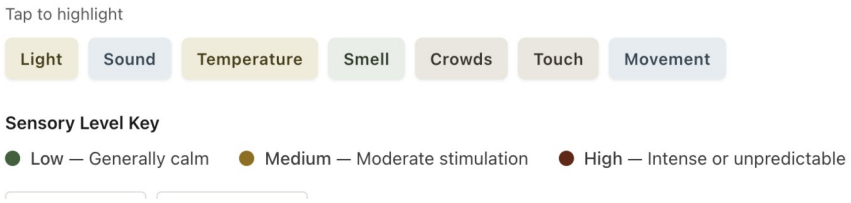

The category filter bar (Light, Sound, Temperature, Smell, Crowds, Touch, Movement). Tapping a category highlights the matching sections. The sensory level key below maps Low, Medium and High to plain descriptions.

Why. Filters serve the *Choice and control* finding, the dominant §3 theme. FG1-P1 used the filter to surface only her own sensitivities; FG1-P5 named the underlying principle. Focus Group 2 sat at opposite ends of the information-density spectrum (FG2-P1: ‘I’m definitely the more information better version’; FG2-P2: ‘the only thing I would want from the sensory information is where the quiet areas are’). The same source content serves both because each reader picks the categories that matter to them.

4.4.2 Public-side: progressive disclosure

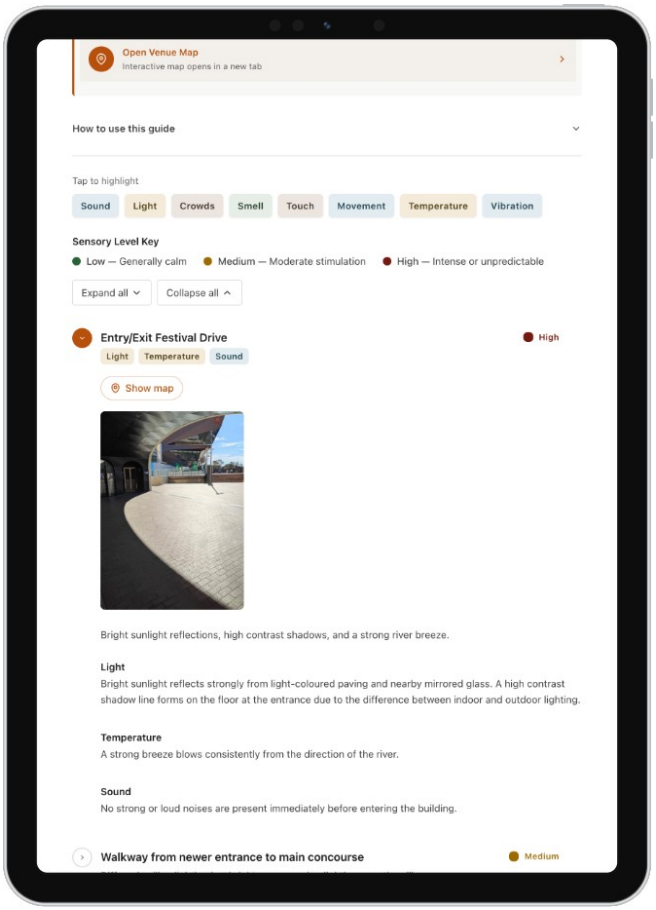

Collapsed area sections with one expanded and the expand-all / collapse-all controls

Why. The design responds to the *Cumulative sensory load* finding. FG1-P1 described the FG1 prototype as ‘too much at once’. The guide now starts simple and expands on the visitor’s terms. FG2-P2 confirmed the collapsed default addressed her initial concern but flagged that scrolling expanded sections still felt heavy: ‘No it’s good here, it’s just when you have to scroll through all of it I don’t like it.’ Sub-section granular control is not yet built (§10).

4.4.3 Public-side: multi-format outputs

The format controls on a published guide: Print (PDF), Save, Word document and Audio narration. The same source content is also available as the interactive web guide and as an embeddable copy for a partner site.

What. A single published guide produces five artefacts: an interactive web guide at a stable public URL, an embeddable version venue operators can drop into their own site, a filterable PDF (with QR codes standing in for inline maps and live links), a Word document, and an audio narration as an MP3 cached per version.

Why. Format choice is not cosmetic on the venue-operator side. Government partners' procurement constraints determine which formats they can publish. The conversation with PTSA reframed publishing as multi-output, with Word and PDF as the immediate distribution path through the operator's accessibility and inclusion resources suite (§6). Visitor-side preferences support the same finding: FG2-P2 signalled an in-app preference, and FG2-P3 described multi-format provision as 'flexibility and options'.

4.4.4 Public-side: auto-generated overview and key facilities

What. Both the top-of-guide overview and the bottom-of-guide key-facilities block are AI-drafted at admin request from the current area content in objective, stimulus-focused language. Nothing references how the environment 'feels'. The admin generates and regenerates each on demand.

Why. The placement responds to a Focus Group 2 hierarchy finding. FG2-P2: 'put everything else that someone might need (eg maps and where quiet spaces are) earlier so they don't bounce off whole website.' Overview at the top and key facilities at the bottom lets a skim-reader see both without committing to a section expansion. FG2-P3 raised in-place editing as a missing affordance: 'what happens if it generates something that isn't accurate, for instance, that would need to be edited.' The admin can regenerate but not edit (§10).

4.4.5 Public-side: maps as optional secondary reference

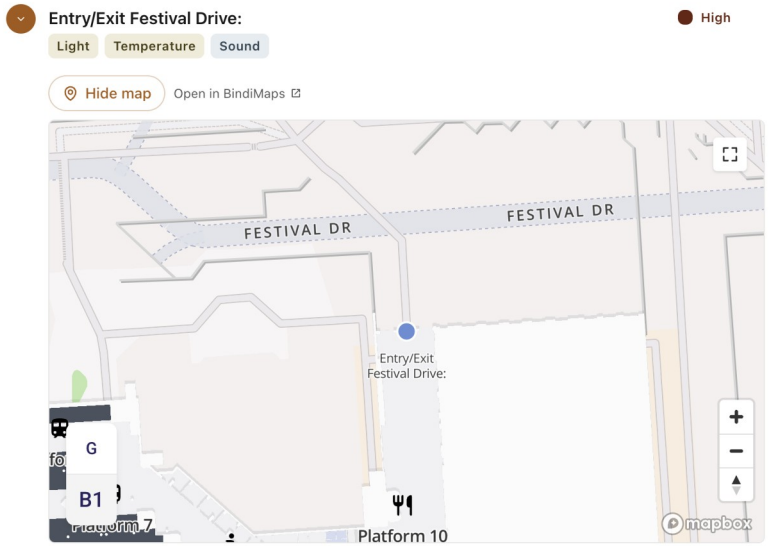

A per-area Show on map mini-map, expanded inside the Entry / Exit Festival Drive area, with the area pinned on the venue map and an Open in BindiMaps link. The map is reachable from each pinned area but is never the entry point to the guide.

Why. The *Photographs over schematic maps* finding cut against an unconditional drop of mapping. Focus Group 1 split between map-anchored and list-anchored entry points. FG1-P5 voiced the map preference: ‘I would rather to start with the map and then click into the detail along the route rather than just reading this list of stuff’. Focus Group 2 praised the per-area *Show on map* mini-map (FG2-P1: ‘Massively, yes’; FG2-P4: ‘a great addition’). Making the map available from every pinned area without making it the entry point anchors map context to specific text and serves both preferences.

4.4.6 Public-side: photographs per section

Why. Photos drew the most converging support across both focus groups. FG1-P1 described seeing photographs in advance as something that ‘takes away a lot of the anticipatory anxiety’. FG2-P1 independently endorsed the photo carousel, and FG2-P4 endorsed the reordering and photo-deletion controls. Focus Group 1 split internally on framing: FG1-P1 leaned toward zoomed-out context shots, FG1-P5 toward entry-direction first-person framing. 360-degree walkthroughs surfaced as a complementary preference alongside still photos (§10). The lightbox and carousel handle photo volume. Framing is left to the auditor.

4.4.7 Public-side: severity and category badges as text

Why. The *Cumulative sensory load* finding showed colour-only coding asked readers to learn a new visual language on top of existing sensory load. Icons compounded that. Text labels remove the decode step. Severity wording itself remains unsettled, with no agreed rubric and Aspect’s preferred terminology pending at time of writing, so the per-venue label override lets each venue use its own vocabulary.

4.4.8 Public-side: responsive design

What. One codebase. Layout, touch targets, typography and navigation adapt from phone to desktop without a separate native app.

Why. Different users prefer different devices, and the same user shifts between phone and desktop across pre-trip planning and in-venue use. A single responsive codebase serves both without the cost of a separate native app.

4.4.9 Admin-side: voice and text data entry on long-form fields

What. Transcription is cloud-backed across devices, with no browser-native fallback, so quality stays constant whether the auditor is on iOS, Android or desktop.

Why. The pattern came from the Focus Group 1 *collaborative authoring* theme. An Aspect facilitator described the workflow:

> *'take a photo, it would again mark it directly in relation to where I am on the map […] upload the photo and upload the detail immediately and then come back and on my laptop just edit the language.' (FG1-P5)*

FG1-P1: 'speak into the app and then it could transcribe it and then automatically put it in. That would be amazing.' On-site auditing in a busy venue makes typing slow and fragile. Voice keeps the auditor's eyes on the environment and hands free for camera and reference materials. Focus Group 2 reinforced the pattern (FG2-P4: 'I like that you can dictate during that section because I know a lot of us like to make voice notes').

4.4.10 Admin-side: in-venue guided wizard

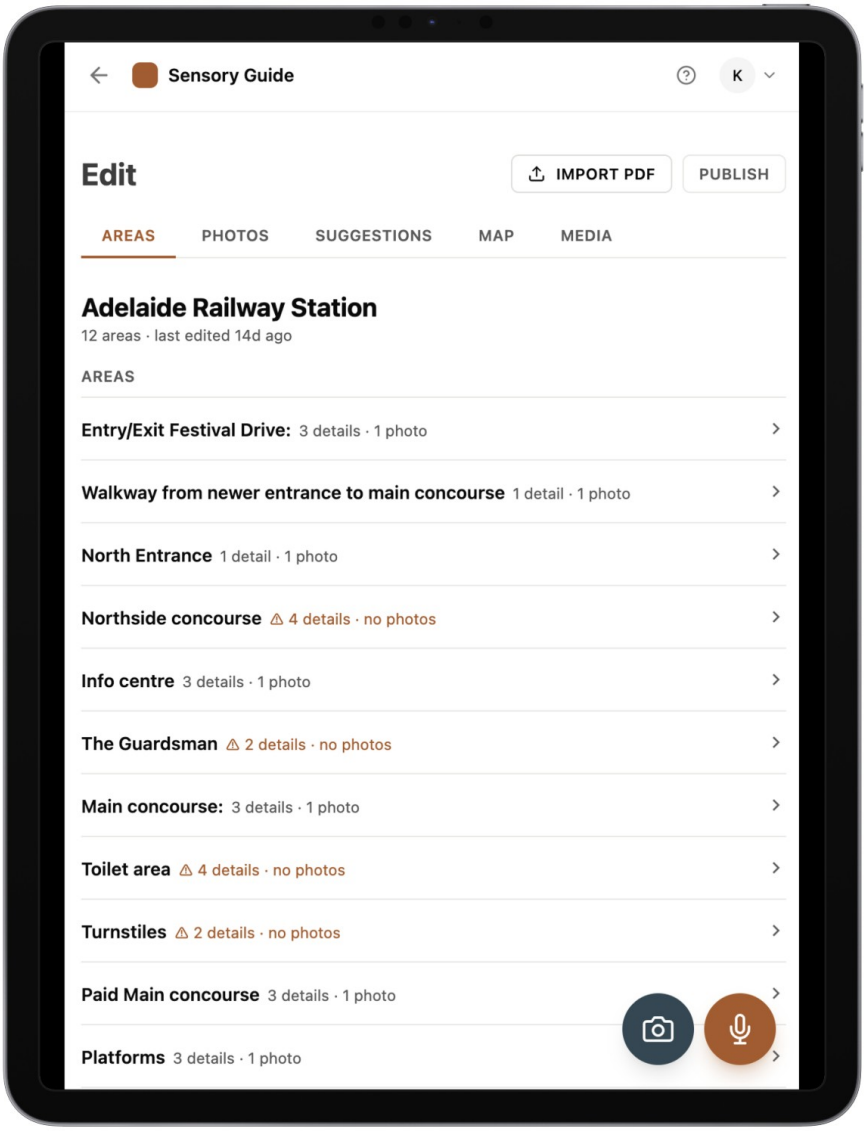

Admin home view of a venue draft (Adelaide Railway Station) with the area list and the two floating quick-capture buttons in the bottom-right corner: camera (new area, jump to photo step) and microphone (new area, jump to summary with voice on). Areas missing photos or details surface inline warnings.

What. Two floating quick-capture buttons on the mobile-entry home (camera and microphone) start a new area at the photo or summary step. Tapping into an area opens a four-step linear flow (photos, summary, sensory details, review) with a tappable progress bar.

Why. On-site auditors need linear, low-cognitive-load capture. The sequence is photos, then a written impression, then sensory categorisation. Focus Group 2 raised that next-buttons were initially hidden. The tappable progress bar replaced the hidden navigation. Focus Group 2 also surfaced a preference among some auditors for entering sensory details before the area summary. Configurable step ordering is not yet built (§10).

4.4.11 Admin-side: AI suggestions with human approval

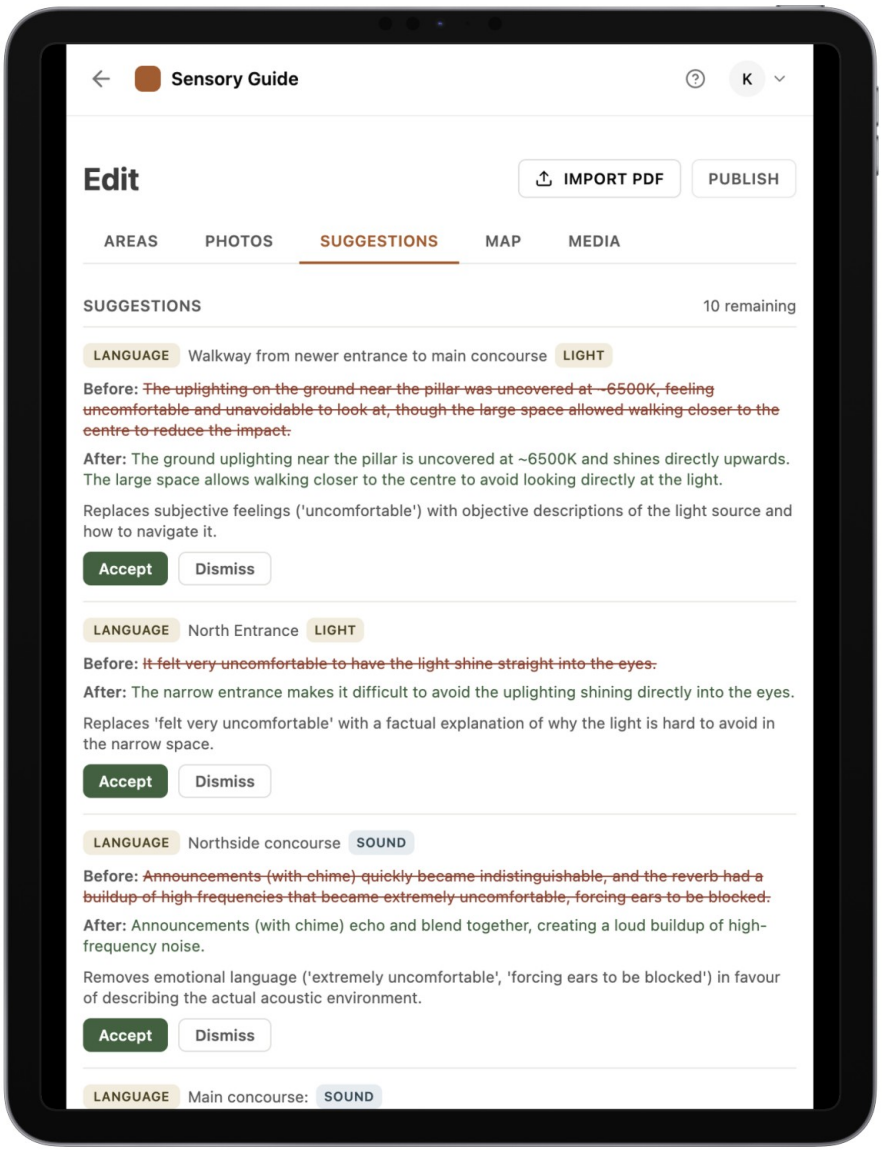

Suggestions tab with three subjective-wording cards stacked, each tagged by area and sensory category. Strike-through Before text sits above the proposed After replacement, followed by a short rationale and Accept / Dismiss controls. The “10 remaining” counter top-right shows the queue size for this call.

What. The panel returns 3 to 15 candidate edits per call, each with a short explanation, original phrase and proposed replacement shown as a strike-through diff. The auditor accepts or dismisses each card. Nothing is silently rewritten. Four interactive suggestion behaviours run from this panel:

1. **Simplify language** (sentence-level). Flag complex vocabulary and propose plainer alternatives.

2. **Remove subjective wording** (sentence-level). Flag emotional or value-laden descriptors against a banned-word list (*unbearable, uncomfortable, overwhelming* and others) and propose objective stimulus-focused replacements. Examples: ‘overwhelming noise’ becomes ‘loud, continuous noise from trains’. ‘Uncomfortable lighting’ becomes ‘bright fluorescent overhead lighting’.
3. **Flag missing sensory details** (section-level). Identify gaps in category coverage and propose which category to add.
4. **Flag cross-area inconsistencies** (document-level). Detect logical contradictions between areas, for example one area rating background music as high while a noisier area rates similar sound as medium.

Three document-level completeness checks (missing venue contact, missing or incomplete address, missing peak-hours information) run once at PDF-import time and appear in a separate post-import panel (catalogue row A22). The auto-generated overview and facility extraction (catalogue row A23) runs on demand and writes to the same objective-language rules as behaviour 2.

Why. The design answers the *AI as helper, not author* finding. Focus Group 1 reacted strongly against the screen-shared AI workflow shown in the admin demo, which ingested an audit PDF and rewrote the content. FG1-P5 asked the AI to ‘check our work rather than changing our work’, and FG1-P1 described the AI-as-error-flagger pattern. AI now proposes. The auditor accepts or rejects each change, and no suggestion is applied without explicit approval. FG1-P3 raised an ethics-of-AI concern, preferring to avoid AI tools personally. The human-approves-every-change rule mitigates but does not resolve that.

Focus Group 2 tested the rebuilt admin tool with three Aspect auditors as prospective admin users, and the AI-anxiety language of Focus Group 1 was absent. FG2-A2, in chat: ‘I really like that you can bypass the AI bits and upload your own words. This is great.’ FG2-P3, accepting a suggestion that flagged ‘hot’ as subjective: ‘It’s sort of like spell check. […] I will accept that.’ The Sensory Guide test with the two lived-experience participants did not exercise the AI suggester directly. Absence of objection is the available signal there, short of direct endorsement. The pattern returns as a candidate framework principle (§9).

4.4.12 Admin-side: multi-admin real-time collaboration

Why. Venue audits in practice involve teams: Aspect for sensory expertise, the venue operator for local knowledge and naming, BindiMaps for content shaping. The need surfaced in Focus Group 1.

FG1-P3: ‘ways to kind of streamline that partner way of working so we can get the experts of, you know, the space, so Adelaide train station team can be inputting all the information for OK, this entry is called this.’

Content also needs ongoing maintenance after publish as venues change, and that maintenance is often a different team member from the original author. Focus Group 2 raised

email-automated editor invites (FG2-A1). The MVP step is generate-link plus copy-message-text (A39), with full email automation flagged for §10.

4.4.13 Admin-side: photo management on mobile

Why. Photos drew the strongest converging support from both focus groups, so management cannot be friction. FG1-P6 raised the absence of a delete affordance during the Focus Group 1 walkthrough: ‘I couldn’t see how to delete any excess photos.’ The current flow covers delete, reorder and cross-area reassignment. FG2-A1, in chat, asked further: ‘can you upload from the photo library, Google Images, Team folder?’ Camera and on-device library are supported. A connection to Google Photos or shared Team folders is not built. Cross-area reassignment partly mitigates by letting editors move misfiled photos without re-uploading (§10).

4.4.14 Admin-side: preview, publish and version history

Why. Venues change. After publish, partners need an audit trail and the ability to revise the guide and publish a new version against the same public URL. FG2-P3 raised a human-readable per-version changelog as missing (‘what changed in this version’). The version-history list exists. A per-version diff does not (§10).

4.4.15 Admin-side: responsive admin UI

What. The same shared draft is the source of truth across phone and desktop.

Why. Capture happens in the venue. Editing happens back at the office. The product treats both as first-class. FG1-P5 described exactly this split: capture on phone, then ‘come back and on my laptop just edit the language or move the photo’. What is captured on the phone is immediately editable on the laptop with no export step in between. The same caveat applies as on the public side: device-to-phase mapping is a useful lens rather than an absolute behavioural prediction.

4.4.16 Admin-side: PDF import as optional starter

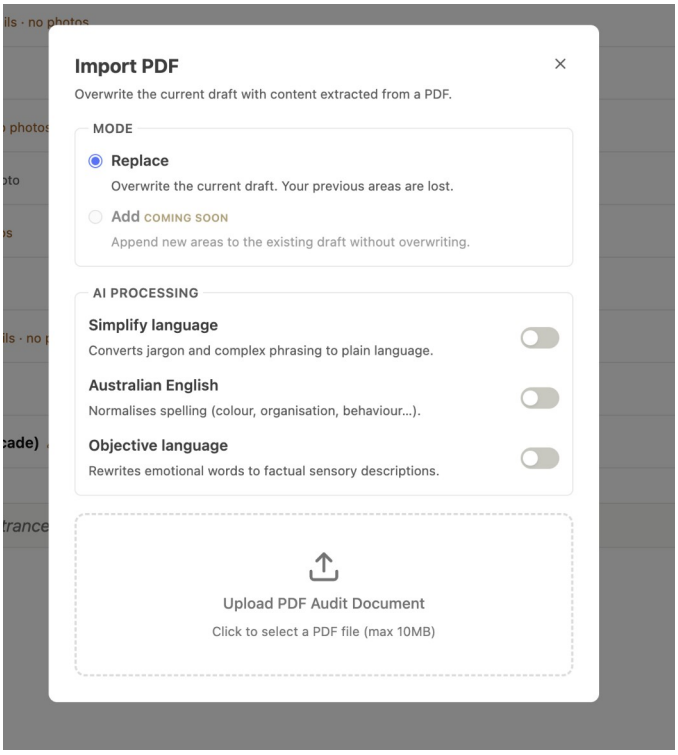

The Import PDF dialog. The three AI-processing options (simplify language, Australian English, objective language) are off by default, so the auditor opts in to each rather than having the import rewrite content automatically.

Why. PDF import was originally positioned as the primary admin workflow. Focus Group 1 relegated it. The lived-experience cohort showed a preference for entering content directly rather than having AI ingest and restructure a finished PDF. Aspect's existing process produces PDFs, however, so skipping the path would force re-typing for that workflow. Imports do not silently rewrite. Every missing item is surfaced to the auditor for explicit action, and that visibility is what allows the AI re-entry path to survive the Focus Group 1 critique.

4.5 What was deliberately not built

Scope discipline is itself a finding.

Excluded	Why not built
Sensory map heatmaps as the primary view	Tested poorly in Milestone 2 and contradicted by Focus Group 1's *Cumulative sensory load* theme. Pivoted to the planning-oriented guide (§3, §5).

Excluded	Why not built
Live transit crowding data	No public Australian transit API exposes occupancy data; Adelaide Metro GTFS-RT has no occupancy-status field (§8).
Real-time crowding alerts and push notifications	Depends on live data that does not exist. No participant raised it as a need.
IoT sound-sensor integration	No public Australian API for venue sound levels; hardware install is out of scope.
Cleaning and maintenance schedules	Internal venue systems with no API access.
Google Popular Times	The Google Places API does not expose Popular Times; scraping violates Google's terms of service.
Runtime data processing on the public viewer	All preprocessing happens at publish time so the public guide stays cheap, fast and cacheable.
Complex wearables and AI prediction systems	Out of scope for the MVP and not raised by participants.
Real-time location detection (turn-by-turn during the visit)	The product is planning-led; live turn-by-turn was not raised as a participant need.
Browser-native speech recognition fallback	Deliberate parity choice; transcription quality variance across native APIs was unacceptable.
Crowd-sourced Waze-style updates	Raised aspirationally by one Focus Group 1 participant; not picked up by others. Post-MVP.
Multilingual content	Not raised in either focus group.
Personal sensitivity profile that persists across venues	Filters persist per venue; cross-venue profile is post-MVP.
360-degree walkthroughs	Wanted alongside still photos in both focus

Excluded	Why not built
	groups (§3 *Photographs over schematic maps*). Post-MVP.

Data-dependency exclusions reflect the absence of Australian public transport and venue data infrastructure (§8). The remainder were scope choices to release a usable MVP. The post-MVP roadmap is set out in §10.

5. Design Iteration

The Sensory Guide reached Focus Group 2 through a build phase before Focus Group 1 and a 56-day iteration window between the two sessions. Five worked cycles trace specific Focus Group 1 findings to changes made before Focus Group 2.

5.1 Method

Development followed a short feedback loop. Research and partner review fed scope choices, small testable changes were made between sessions, and each focus group tested the cumulative work. Two elements defined the method:

- **Structured planning per feature.** Each major capability passed through a planning chain (product brief, requirements, architecture sketch, epics, user stories) before any code was written. The chain produced auditable artefacts that other team members and the partner organisations could read.
- **Per-feature iteration cycles.** Each feature moved through a short cycle of build, internal review and partner review. Aspect's Autism Friendly team reviewed iteration outputs between the two focus groups, supplying lived-experience feedback on smaller changes without waiting for a formal session. Automated quality checks ran on every commit.

5.2 Pre-Focus-Group 1 build

A Milestone 2 prototype (December 2025) tested the action-plan format as a popup over the live map, with sensory categories as colour-coded polygon overlays and a print-friendly layout, from a hardcoded example dataset.

Milestone 3 began with a greenfield rebuild to add multi-venue authoring for auditors and venue editors, persistent storage, and content that could scale beyond a single hand-built example.

Development on the new codebase began on 24 January 2026, just under nine weeks before Focus Group 1. The Focus Group 1 prototype was a live working Sensory Guide for Adelaide Railway Station alongside a screen-shared admin PDF-ingest demo, in which auditors uploaded an Aspect audit PDF and the AI generated the structured guide for review.

5.3 Iteration breadth: Focus Group 1 to Focus Group 2

Fifty-six days separated Focus Group 1 (24 March 2026) and Focus Group 2 (19 May 2026). Iteration across the window covered AI authoring behaviour, mobile capture, admin workflow, the in-app map view and area pinning, filter behaviour, severity terminology and accessibility, each driven by a Focus Group 1 finding.

The chart groups 234 code changes across thirteen themes, with admin-workflow restructure and the in-app map the largest areas. Each bar counts individual code changes, which include the automated story-by-story changes described at §5.1, so the figures show where effort concentrated rather than a tally of separate decisions. A further 37 housekeeping and documentation changes sit outside the thirteen themes.

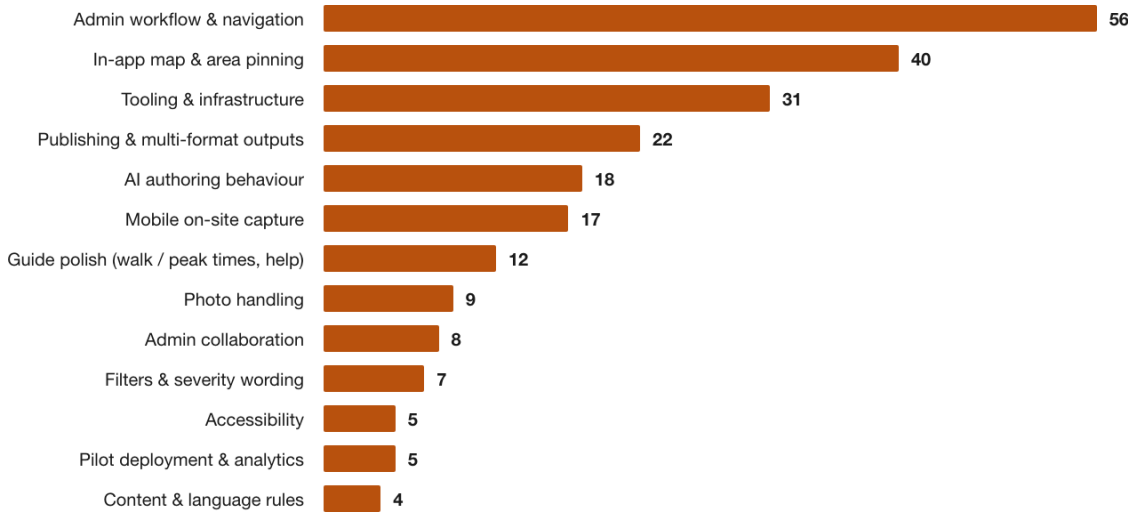

5.4 Iteration depth: five worked cycles

5.4.1 AI as author to AI as suggester

The original admin tool was built around a single workflow: an Aspect auditor uploaded an audit PDF, the AI ingested and rewrote the content as a structured guide, and the auditor reviewed the AI's output at the end. The premise was minimal disruption to Aspect's existing workflow, with the AI as author and the auditor as proof-reader.

Focus Group 1 rejected it. FG1-P5 raised non-determinism as a trust issue, ‘every time you upload the same document, it would work differently’, and named the inversion the group wanted: ‘check our work rather than changing our work and us checking it.’

The admin tool was rebuilt to invert the relationship. Auditors now fill in venue summary, area name, sensory categories, photos, severity and prose directly, on a phone or laptop. The AI was moved to a separate Suggestions panel that scans the draft and proposes sentence-level edits with a rationale. Each suggestion is individually accepted or rejected. Nothing is auto-applied. PDF import remains as one of several optional entry paths rather than the default.

The Suggestions panel catches things human review alone misses: subjective wording slipping into descriptions (FG1-P3, an Aspect auditor, named the pattern in Focus Group 1 with examples like ‘noticeable temperature’, ‘comfortable’, ‘unbearable smell of ammonia’), and missing context where an area entry lacks a category, photo or severity.

Focus Group 2 tested the rebuilt tool with Aspect’s Autism Friendly team. FG2-P3, accepting an AI suggestion that flagged ‘hot’ as subjective:

> *‘It’s sort of like spell check. I would go, oh, actually that is, I did spell that wrong. I could have done better. I will accept that.’*
>
> *(FG2-P3, Focus Group 2, 19 May 2026)*

5.4.2 Filter toggles to auto-expanding filters

In the Focus Group 1 prototype, filters highlighted matching tags only inside areas the visitor had already manually expanded. A visitor filtering for ‘touch’ saw the active-count change but had no way to find which areas contained the matching content without opening every collapsible section.

FG1-P2 described the behaviour:

> *‘some of the tags didn’t directly show up […] I would have to open each individual tab […] to make sure that I had got all of the information.’*
>
> *(FG1-P2, Focus Group 1, 24 March 2026)*

The filter was changed to auto-expand any area containing a matching category, and to collapse sections again when filters are cleared. A returning visitor’s last open-or-closed state is preserved across page loads.

Focus Group 2 used the feature live during the walkthrough without targeted comment.

5.4.3 Manual photo workflow to on-site mobile capture

Focus Group 1 surfaced auditors’ demand for on-site capture. FG1-P5 described what they wanted:

> *'take a photo, it would again mark it directly in relation to where I am on the map […] upload the photo and upload the detail immediately'*
>
> *(FG1-P5, Focus Group 1, 24 March 2026)*

A mobile-first capture flow was built into the admin app. A phone or tablet running the web app can add an area, take a photo with the device camera, dictate the description, and assign the area to a category, all on-site, with no document-assembly step on a laptop afterwards.

Two large quick-add buttons on the area list jump straight into the camera or the microphone. Voice dictation auto-recases the transcript (title case for names, sentence case for descriptions).

Focus Group 2 endorsed the change. FG2-P4: 'I like that you can dictate during that section because I know a lot of us like to make voice notes on our phone.' The mobile-capture flow was used through the walkthrough without complaint.

5.4.4 Icons on a map to text categories with map on demand

The icon-overlay-on-map design was retired before Focus Group 1. The public guide now leads with a text-led action plan: areas as collapsible sections with category badges, level pills and prose descriptions. Categories are colour-coded inline next to the text they describe.

A small per-area 'Show map' button reveals an inline map view zoomed to that area only, on demand. The sensory level key was lifted out of the page footer to sit immediately below the filter bar. Category colours were tuned for WCAG 2.2 AA contrast and to keep adjacent severity levels visually distinct.

FG1-P1 described the earlier map-led prototype as too much to parse at once (§3).

Focus Group 2 endorsed the per-area on-demand map. An Aspect facilitator asked the room whether others found the per-area map helpful; FG2-P1: 'Massively, yes.'

FG2-P1 also proposed extending the inline category colour-coding from the badges into the prose itself: 'I just feel like that would be an easy way for my brain to go, oh, that means that, that means that, that means that, without having to actually read it really in depth.' The suggestion sits in the post-MVP roadmap (§10).

5.4.5 Fixed severity wording to configurable labels

Focus Group 1 flagged the default severity wording ('Low / Medium / High' paired with 'generally calm / moderate stimulation / intense or unpredictable') as too thin and partly subjective. FG1-P5:

> *'I also just think that generally calm, moderate stimulation, intense or unpredictable isn't enough detail and is a bit subjective. We do have […] that we've developed about why we rate things as low, medium, high. So I would just incorporate that in there.'*

> *(FG1-P5, Focus Group 1, 24 March 2026)*

Severity was split into two layers. The underlying values (low, medium, high) stayed fixed so data and badges remained consistent across venues. The displayed labels became per-venue configurable: a venue can override the wording of low, medium and high to match its auditor's house terminology.

Defaults remain 'Low / Medium / High' when no override is set. Earlier guides keep working unchanged. Aspect can now adopt preferred wording across every Aspect-authored guide without a code release.

The platform-side capability is in place; the content-side wording is pending Aspect's preferred terminology. FG2-P3 repeated the severity-wording concern in Focus Group 2 (§3).

5.5 How risks were managed in iteration

Iteration managed three project risks.

- **Schedule and delivery risk.** Per-feature iteration with milestone buffers held the schedule. Focus Group 2 ran on its planned date (19 May 2026) without slippage.
- **Scope-creep risk.** Candidate features were folded into existing cycles or deferred to the post-MVP roadmap.
- **Partner-availability risk.** Aspect reviewed iteration outputs between focus groups, and platform-side capability was released ahead of partner content where partner wording was the blocker.

6. Pilot Deployment

6.1 Pilot site and partnership

The Milestone 3 pilot ran at Adelaide Railway Station, the underground rail hub at Festival Drive and North Terrace on the northern edge of the Adelaide CBD. The station is an existing BindiMaps customer, so indoor map data was already published for the venue and the pilot did not need to build the underlying map first.

Public Transport South Australia (PTSA), the customer-facing public-transport authority under the South Australian Government's Department for Infrastructure and Transport, is the pilot's operator partner. The rail network it operates serves Adelaide Railway Station; the wider transit network it manages is Adelaide Metro. The pilot's partner contact at PTSA was its Customer Strategy and Inclusion Lead, who corresponded by email across the pilot.

Three parties shaped what the pilot could test. PTSA brought operator-side content authority and procurement constraints. Aspect (§3) ran the on-site sensory audit and the lived-experience review. The BindiMaps platform supplied the indoor map data, hosting and the publishing pipeline.

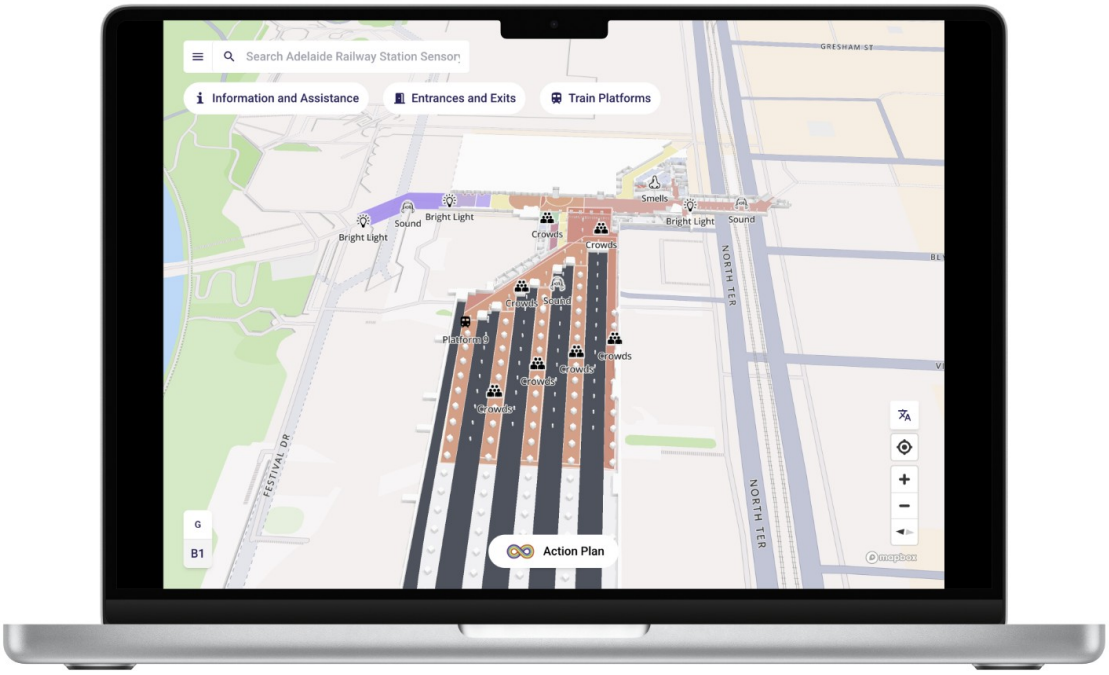

Adelaide Railway Station context: Festival Drive / North Terrace entry and the south-side concourse covered by the Aspect audit

6.2 Deployment

The pilot guide was authored from the Aspect on-site audit and published through the admin workflow (§4.2). The static layer carried nine sensory categories (sound, light, crowds, smell, touch, movement, temperature, vibration, air quality) with severity badges and per-category descriptions for each area along the audited path. The public guide is live at a stable URL, https://sensory-guide.web.app/venue/adelaide-railway-station, on the Sensory Guide platform at https://sensory-guide.web.app/ (authoring and editing require an operator login). A single publish action regenerated every downloadable artefact.

Five artefacts published from a single source draft.

- **Interactive web guide** at the public URL, with filters, progressive disclosure, photographs and per-area *Show on map* mini-maps.
- **Operator-hosted version** of the live guide, dropped into the operator's own webpage and updating from the same source.

- **PDF** with QR codes substituting for the interactive map and per-area embedded media, filterable so a visitor can download the full guide or only the sections matching their active filters.
- **Word document** export of the same content.
- **Audio** narration of the guide as an MP3, generated from the published text.

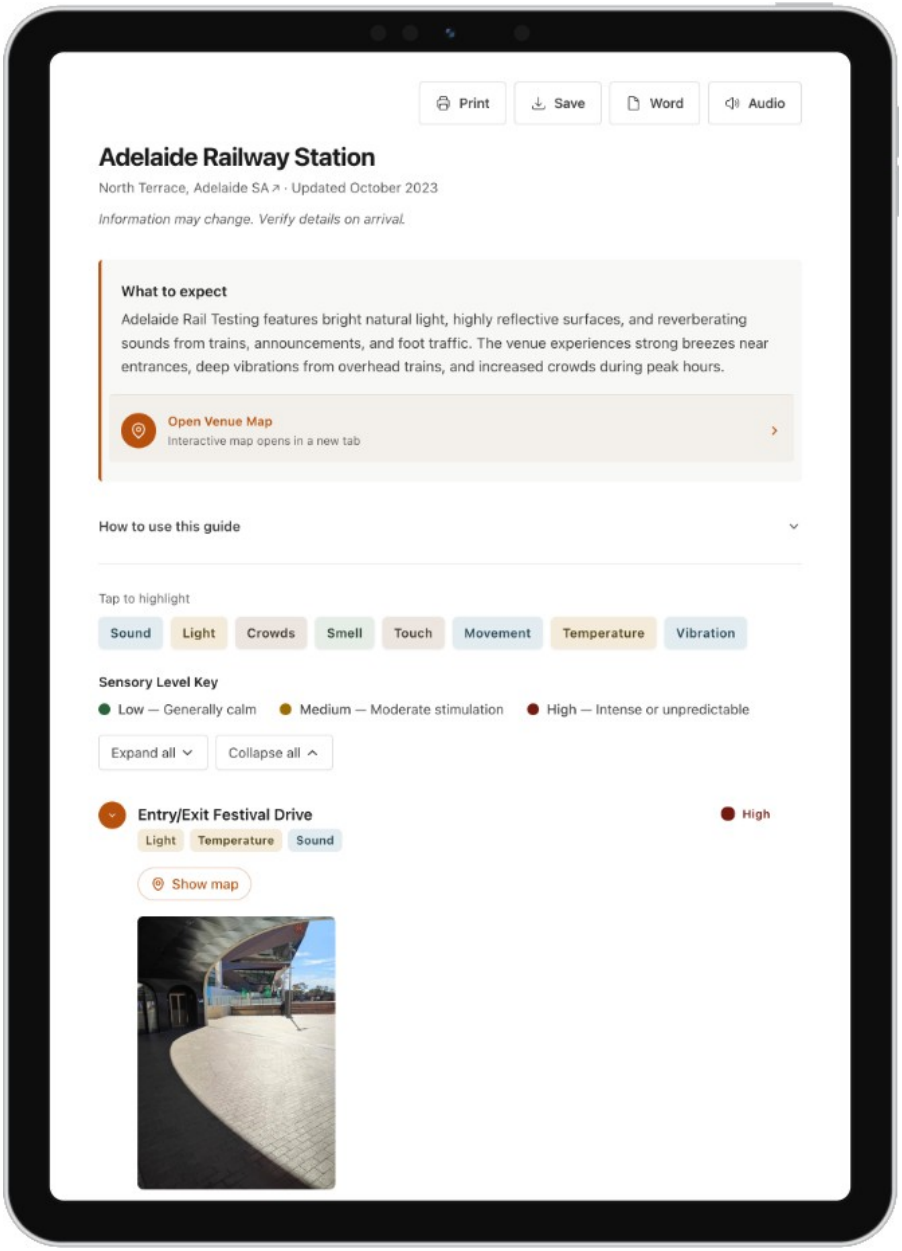

Pilot version of the Adelaide Railway Station Sensory Guide at its public URL

The original distribution plan was to host the live interactive guide directly within the Adelaide Metro website, on the assumption that the operator's own site is where visitors look first for accessibility information.

PTSA's procurement framework did not accommodate third-party-hosted interactive content on government infrastructure inside the pilot window. PTSA's account was that the live interactive version would have required a security review with no clear timeline. The format PTSA could distribute without that review was a document uploaded to their accessibility and inclusion resources suite.

Publishing was reframed for the rest of the project. Web became one of several equal-weight artefacts authored from the same draft (§4.4.3, §9.1.4).

Each format reshapes the content at publish time. The PDF inserts a QR code in place of the inline map. The audio drops visual references and is narrated for an audio-first listener. The Word document carries plain content suitable for transfer into a third-party template.

The PDF, Word and audio artefacts lose the interactivity and live-update path the web copy preserves. QR codes route readers from print back to the live web copy when they are online.

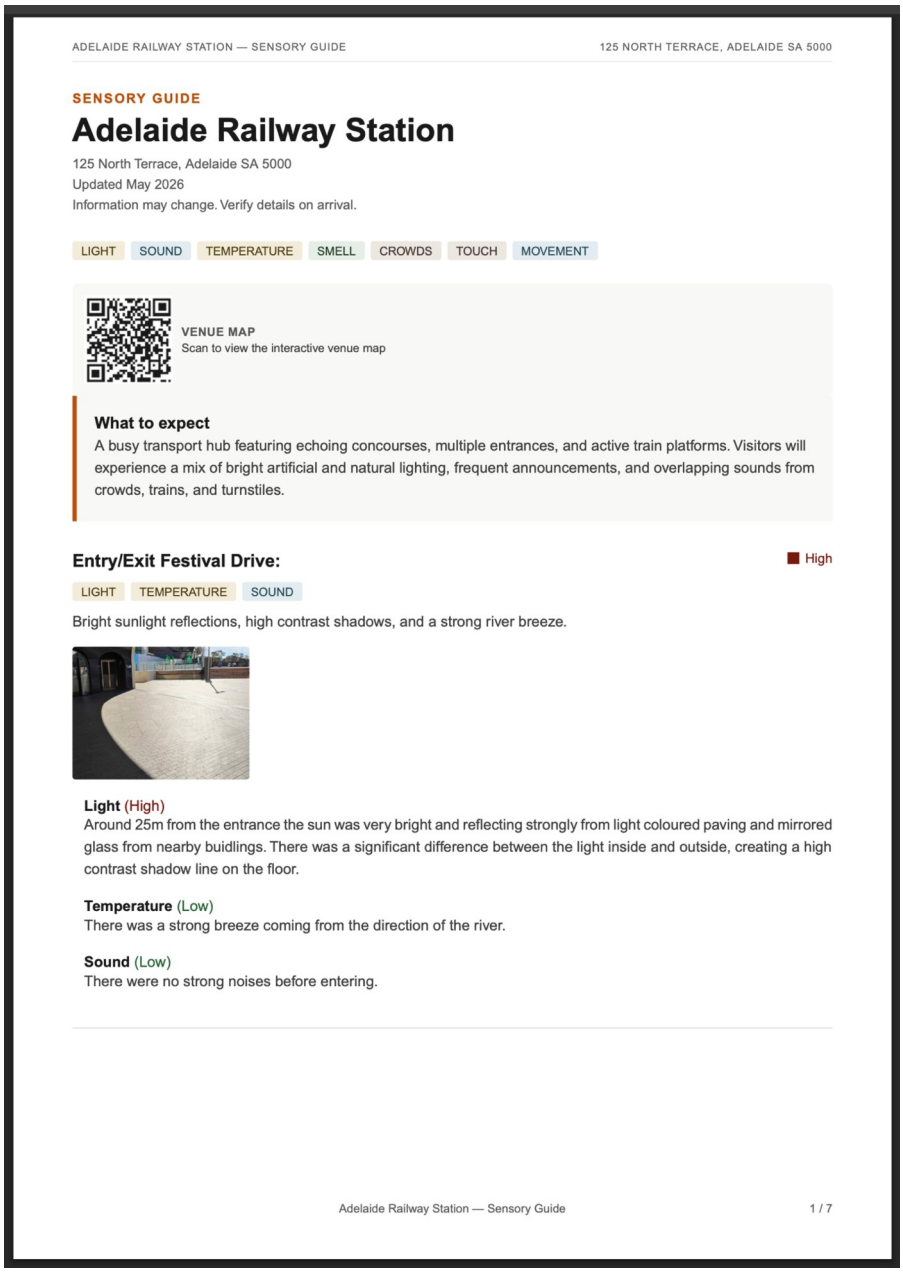

PDF export of the published Sensory Guide, one of the formats delivered to PTSA. The QR code replaces the inline venue map, linking back to the live interactive copy.

6.3 Operator-supplied dynamic content

Live transit-crowding data was unavailable for the Adelaide Metro network within the pilot window. With no live feed available, the changing information came from the operator instead: the guide is revised and re-published through the admin tool as the venue changes.

For the in-moment user-facing features, operator-authored static content stood in. Sensory category filters and per-area severity badges serve as the planning-time analogue of in-moment

alerts, with matching sections auto-expanding on filter selection (§4.4.1, §5.4.2). Journey-ordered area sections and per-area *Show on map* mini-maps serve as the route preview (§4.3.1, §4.4.5).

After the audit, PTSA flagged a turnstile-to-smart-gate replacement at the station. The admin tool's re-publish workflow updated the affected entries; partner photographs were pending at time of writing.

The version-history view (§4.4.14) records every published snapshot and supports revert. A renamed venue is flagged as unpublished until it is re-published. Audit content decays as venues change, so currency depends on the admin pipeline rather than the static audit (§9.1.2, §10). Aspect's twice-yearly Quiet Room re-audit cadence is the standing precedent for on-cycle revisions.

Time-of-day advice authored by the auditor (typical peak hours, typical cleaning windows) gives the predictability a live feed would otherwise have surfaced. The user-research cohort leaned toward preparation before travel over alerts in the moment, and toward journey-ordered planning over in-moment route guidance (§3, §8).

Aspect editors in Focus Group 2 surfaced a missing email-automation step in the editor-invite flow, which currently relies on a link and a copy-message action (§10).

6.4 Partner experience

PTSA's Customer Strategy and Inclusion Lead wrote back to BindiMaps after receiving the published guide:

> *'This sensory guide seems like a very helpful resource which we are excited to add to our customer accessibility resources!'*
>
> *(Customer Strategy and Inclusion Lead, PTSA, email correspondence)*

On the partner side, the PDF was added to PTSA's accessibility and inclusion resources suite, which PTSA shares with its stakeholders.

The partner contact named the Adelaide Metro website as the longer-term distribution path, once PTSA's Digital team has set up a suitable location for resources of this kind.

The same contact also reported a recent internal PTSA catch-up on the BindiMaps roadmap, which raised both the separate Sensory Guide and an in-navigation sensory overlay as possibilities.

Two limits apply. Operator-side feedback came predominantly through one partner contact, so broader PTSA organisational uptake is a post-pilot question. The roadmap conversation is a forward signal rather than a delivered outcome. Both are taken up in §10.

7. Evaluation Metrics & Results

7.1 Methods and instruments

Four sources fed the evaluation.

- **Focus-group transcripts.** Two Milestone 3 sessions: Focus Group 1 (24 March 2026) testing the sensory-map prototype with four lived-experience participants; Focus Group 2 (19 May 2026) testing the redesigned Sensory Guide with two Autistic participants who are not Aspect staff (FG2-P1 and FG2-P2) and the admin tool with three Aspect auditors. Quotes transcribed verbatim and thematically coded (§3).
- **In-product analytics.** Every interaction on the public guide and the admin tool is logged through a typed event catalogue, feeding Google Analytics on the public side and an equivalent analytics service on the admin side (P31, §4). Microsoft Clarity adds session replay and heatmaps. A super-admin aggregate dashboard shows these across all venues (A43).
- **In-product feedback.** The public guide carries an in-page thumbs-up / thumbs-down with a free-text comment up to 100 characters (P22); the admin tool exposes a per-venue feedback view with 1, 2, 4-week and all-time filters (A40).
- **Audits.** A WCAG 2.2 AA accessibility audit and a security audit were performed on the build on 19 May 2026, each combining automated scanning with manual inspection. Automated accessibility, performance and dependency security checks (Pa11y, Lighthouse, dependency audits) also run against every code change as part of the deployment pipeline.

The Milestone 3 sessions used qualitative methods. No pre / post planning-anxiety and confidence instrument was administered. The stress and confidence change reported in §7.3 comes from participant accounts of planning with and without the artefact. Items not measurable inside the Milestone 3 window (a standardised pre / post instrument, independent penetration testing, third-party WCAG sign-off, visitor-scale engagement capture, live-data alerts) are named in §10 as broader-deployment additions.

7.2 Engagement

Per-event tracking covers the interaction points the focus-group findings touched: section expansions (the progressive-disclosure sections, §4.4.2, §5.4.2), filter toggles by category and severity (§4.4.1), per-area *Show on map* opens (§4.4.5), image views, external-link clicks, PDF and audio download events, and feedback submissions. Microsoft Clarity provides session replay for behavioural review.

The super-admin dashboard (A43) shows aggregate counts across all published venues. Per-event counts are deferred to the post-pilot capture, since the Milestone 3 distribution window was partner-internal (§6.3).

7.3 Stress and confidence change

The change in user-reported stress and confidence is reported qualitatively.

Focus Group 1's reception of the original sensory-map artefact (§3) was consistently one of overwhelm, hierarchy confusion and anticipatory anxiety. Focus Group 2 reception of the redesigned guide was the converse: per-area mapping, category structure and progressive disclosure matched participants' planning behaviour, and the overwhelm language did not reappear. FG2 reception quotes are at §7.4.

Focus Group 2 did not re-test the original sensory map, so the comparison is the participants' reception of the redesign, not a controlled before-and-after on the same cohort.

The paired quotes below set the two cohorts' reception side by side.

Focus Group 1, sensory-map prototype (overwhelm)	Focus Group 2, redesigned Sensory Guide (confidence)
'I did struggle to understand all of the information, and I think in part it was because there was competing things going on. So my pattern recognition was kind of jumping between things and not quite sure on like the hierarchy of what to start with first and follow to.' (FG1-P1)	'When I was a kid and I would catch a bus for the first time somewhere, I would have to know every bus route, every timetable. […] it reduces my anxiety and makes me feel better about being out in public because it reduces the layer of unpredictability.' (FG2-P1)
'I thought the information was great, it was just kind of too much at once and then I was overwhelmed and then I couldn't really like piece it through.' (FG1-P1)	'I like the addition of sound, light, temperature. […] I think that's a really good way of measuring things.' (FG2-P1)
'trying to work that all out at the same time with maybe not enough context of what different colours indicate or different intensities indicate […] makes it really complex.' (FG1-P3)	'there's a bit of a warning here. There's some things missing, so there's no photos. I love that as well because it reminds me of what I need to still do.' (FG2-P3)
'a sensory map currently could be extra confusing because there's a lot of visual information. Some of it is a new language to	'I really like that you can bypass the AI bits and upload your own words. This is great.'

Focus Group 1, sensory-map prototype (overwhelm)	Focus Group 2, redesigned Sensory Guide (confidence)
learn.’ (FG1-P5)	(FG2-A2, chat)

7.4 Satisfaction

Satisfaction is reported qualitatively from the two focus groups and the PTSA partner correspondence (§6.4). No Likert-scale survey was administered. The in-product thumbs-up / thumbs-down with comment (P22) is built but did not accumulate a public visitor sample inside the partner-internal pilot window (§7.2).

The quotes below are positive-reception evidence. Ambivalent and critical material, including the FG2-P1-and-FG2-P2 information-density disagreement and the bug reports raised in Focus Group 2, sits in §3.4.

Reception of the redesigned Sensory Guide across both Milestone 3 sessions was positive within the constraints of the partner-internal pilot window. The lived-experience cohort responded positively in session (FG2 quotes in §7.3); the Aspect auditors confirmed the rebuilt admin tool works for their workflow (FG2-P3 received the publish-readiness warnings positively, quoted at §9.2 G20); and PTSA’s primary contact endorsed the artefact for inclusion in the operator’s accessibility resources suite (§6.4).

PTSA’s partner satisfaction is recorded in the Customer Strategy and Inclusion Lead’s first email on receiving the published guide, welcoming it as a resource to add to the operator’s customer-accessibility set (quoted at §6.4).

7.5 Quality metrics

The build was reviewed against WCAG 2.2 AA for accessibility and against a standard security checklist. The accessibility review found no critical or high-severity issues. The security review flagged two high-severity issues, both fixable in under 30 lines of code per the auditor. Automated checks of the same kind run against every code change as part of the deployment pipeline.

Security

The 19 May 2026 security review, combining automated scanning with manual inspection, covered database and file-storage access rules, server function authorisation, client-side injection risks, secrets handling and third-party dependency security. The auditor approved continued deployment on condition the two high-severity issues are remediated.

Beyond those, the review noted four medium-severity issues and four low-severity issues. Database and file-storage access rules, client-side injection paths, secrets handling and cross-origin restrictions returned no issues.

The full severity table and the auditor’s note appear at Appendix B.1; the full audit is on file with the project team.

Accessibility

The 19 May 2026 accessibility audit, combining automated scanning with manual inspection against the WCAG 2.2 Quick Reference, covered the public guide, the admin tool, the audio output and the design-system colour tokens. No issues were rated critical or high.

Five medium-severity findings were noted, all in colour contrast or motion, each with a planned mitigation. The audit also flagged eight low-severity best-practice items and two informational observations. The auditor confirmed that none of the findings block release.

Page structure, native interactive elements, visible keyboard focus indicators across the public guide, screen-reader cues on expandable sections and filter buttons, live regions on error and save-status messages, and large touch targets on the audio button were all clean.

The full severity table and the auditor’s note appear at Appendix B.2; the full audit is on file with the project team.

Automated checks on every code change

Automated checks run on every code change and block deployment if any fail.

Each change is run against code-style linting, unit tests, a third-party dependency security audit, Lighthouse performance and accessibility scoring against the production build (three runs, median taken; blocking thresholds accessibility 0.95 and performance 0.80), and Pa11y accessibility scanning of the public homepage against the WCAG 2.2 AA profile. A change with any failing check cannot reach production.

Pa11y coverage is currently the public homepage only. Admin pages and individual venue guide pages are covered by the manual audit above but not by automated checks on every change. The coverage extension is a §10 item.

On the most recent passing build (27 May 2026, default mobile profile), Lighthouse returned accessibility 100 and performance 92, both clearing the blocking thresholds, with best-practices 79 and SEO 90 at advisory level. The Pa11y homepage scan passed with zero errors.

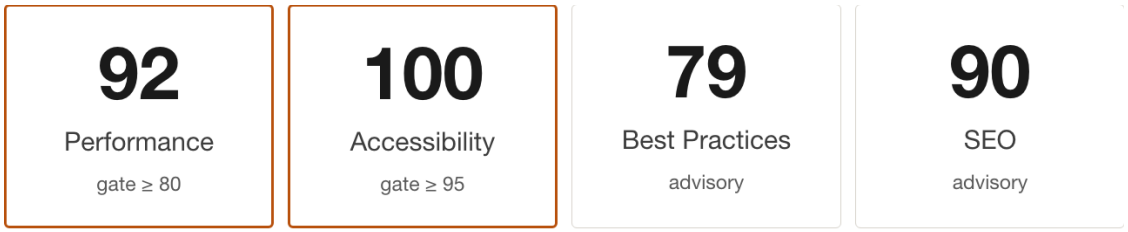

PA11Y WCAG 2.2 AA scan of the public homepage: passed, zero errors.

Bundle size

The public viewer downloads approximately 144 KB on first load (compressed), measured against the build of 27 May 2026. Most of this is the application's own code (116 KB), with the React library and the stylesheet making up the rest. The sign-in and editing-database code is not part of this download. It loads only when an operator signs in to the admin tool, or when a public visitor asks for a guide's spoken-audio version. The other large admin resources behave the same way and stay off the public path: PDF generation is 512 KB and Word-document generation 101 KB.

What this means for a visitor depends on their connection. At 144 KB compressed, the initial download transfers in roughly three-quarters of a second on the throttled mobile connection the audits simulate (about 1.6 Mbps), and in a fraction of that on a typical unthrottled 4G connection. Time to a fully interactive page adds the cost of parsing and running the code on the device, which the production-build measurement noted in §10.4 will capture; the figures here describe the download itself.

Total uncompressed application size across all chunks is 3.5 MB. Full per-chunk breakdown in Appendix B.3.

7.6 Limitations

The evidence above rests on a small cohort: six lived-experience participants across the two focus groups plus three Aspect auditors on the admin tool, with Aspect's recruitment channel concentrating coverage on Autistic experience (§3.5).

The pilot covers one venue, with no pre / post numeric instrument administered. Because the guide was shared with the partner rather than launched to the public during the pilot window, too few real visitors used it to support engagement claims at any scale. The audits are automated and do not substitute for independent penetration testing or real-user assistive-technology session work.

Live-data alerts (transit-crowding, sound levels) that would support in-moment in-product alerting were not measurable in the Milestone 3 window because the dynamic data feeds were not available for the Adelaide Metro network (§8).

Inside the evidence base: direction of effect, feature-by-feature user reception, two passing audits (accessibility unconditional, security conditional) and automated-check enforcement of the engineering thresholds on every code change. Outside it, and named in §10: population-scale visitor engagement, quantified pre / post anxiety changes and externally certified audit attestation.

8. Data-Dependency Findings & Pivots

The project began with a set of intended features that assumed several live external data feeds. Most of those feeds were not available for the Adelaide Metro network or comparable Australian transit hubs.

8.1 Originally intended dynamic data sources

Four dynamic inputs were named at project scoping. A fifth (Google Popular Times) emerged during feasibility work as a substitute for the missing transit-crowding feed.

- **Live transit crowding data.** Real-time occupancy on inbound and outbound services serving the station, on the model of Transport for NSW's almost-live train-crowding feed.
- **IoT sensor data, particularly sound levels.** Venue-side monitoring of ambient sound at concourse and platform level, surfaced into the guide as a live reading.
- **Maintenance and cleaning schedules.** Operator-side facility schedules surfaced as a live or near-live signal of when a given area was likely to be louder, busier or wet underfoot.
- **Real-time alerts.** A composite visitor-facing feature combining the three feeds above, presenting condition warnings as they unfolded.
- **Google Popular Times.** The live and historical busyness data shown on Google Maps listings, considered as a proxy for venue-level crowding where transit-side occupancy was absent.

8.2 Feasibility findings per source

Each source was tested against published APIs, terms of service and the platform's operational scope.

Intended source	Outcome	Reason
Live transit crowding (Adelaide Metro)	Infeasible	Adelaide Metro GTFS-RT does not publish the

Intended source	Outcome	Reason
		`OccupancyStatus` field
IoT sound-level sensors at Adelaide Railway Station	Infeasible	No public API; would require hardware deployment outside the platform's scope
Maintenance and cleaning schedules	Infeasible	Internal operational systems with no public API
Real-time alerts (composite)	Infeasible	Has no underlying feed once the three above are subtracted
Google Popular Times (substitute)	Infeasible	Not exposed in the Google Places API; scraping Google Maps breaches its terms of service

Live transit crowding (Adelaide Metro)

The GTFS-RT specification defines an `OccupancyStatus` field that any operator can populate. The gap is in publishing rather than protocol design.

Adelaide Metro publishes a static GTFS feed and a SIRI feed for real-time arrivals, but its GTFS-RT feed does not populate `OccupancyStatus`. Transport for NSW does publish crowding data on its data portal, but the feed covers NSW services only. No other Australian operator publishes occupancy for the Adelaide network.

The MVP therefore has no live-occupancy signal available.

IoT sound-level sensors at Adelaide Railway Station

No public API exposes ambient sound-level data for Adelaide Railway Station or comparable Australian transit hubs. Sourcing it would have required installing hardware sensors at the venue and building a bespoke data integration. Both fell outside the platform's scope and outside BindiMaps' operational control over a third-party site.

Maintenance and cleaning schedules

Maintenance and cleaning schedules at major transit hubs sit in internal operational systems with no public API at the time of the pilot. Surfacing them would have required a per-venue data-sharing arrangement with the operator's facilities function.

Real-time alerts

Real-time alerts at Adelaide Railway Station were originally scoped as a composite feature drawing on the three feeds above. With each of those feeds infeasible, no live alerts feature was built.

The underlying intent was to warn a visitor of varying conditions. The MVP carries the equivalent information statically. Peak-hour patterns, typical cleaning windows and similar predictable variation are authored into the Sensory Guide content by Aspect auditors for the visitor to plan around.

Google Popular Times

The Google Places API (the data product) does not expose Popular Times or busyness data. Third-party services such as BestTime.app reach the data by scraping Google Maps (the consumer surface) and resell access for roughly $60/yr. That access rests on a terms-of-service breach the platform cannot control and could be cut off without notice, so it is not a dependency a published guide could safely rely on.

All three available paths (direct scraping, paid third-party scraper, waiting for an official Places API field) either breach the Google Maps terms of service or have no public signal of arrival. The MVP does not use Popular Times.

8.3 Convergence with the user research

The data-feasibility findings alone would justify the pivot to operator-supplied static content. User research in §3 points the same way on one specific point. Live alerts did not surface as a primary need across either Milestone 3 focus group. Focus Group 1 participants and FG2-P1 described value in preparation before travel (FG1-P1 and FG2-P1 quoted at depth in §3.3).

Technical infeasibility and user direction converge on the alerts feature. The underlying feeds are not available for Adelaide Railway Station, and the planning-dominant pattern in the cohort favoured preparation in advance over warnings as conditions unfolded.

The convergence is narrower than ‘static beats live in general’. FG2-P2 described in-venue lookup of static content on arrival (‘more likely when I arrive, unless the venue has told me this exists’). That use pattern needs no live alerts at all, and the MVP supports it.

Preparation-oriented, audit-authored content is the design that best matches the alerts-axis evidence. Live signal retains value for other features and other cohorts.

8.4 What was delivered as dynamic content

With the live data feeds unavailable, dynamic content comes instead from venue staff editing the guide through the admin tool. Auditors and venue editors revise the draft and re-publish the

live copy, with typical peak hours and cleaning windows authored into the area entries (§4.2, §4.4.3).

Updates happen on the cycle of venue change rather than second-by-second. The §3 cohort accepted this trade-off, given its planning-dominant pattern. Transferability to other user groups is open.

8.5 Lessons for transferability

Two lessons follow for venues and partners that adopt the framework beyond Adelaide Railway Station.

- **Australian transit operators that do not currently publish OccupancyStatus should publish it.** The field is part of the GTFS-RT protocol and Transport for NSW already populates it. Adelaide Metro does not (§8.2). Closing the publishing gap on the Adelaide network, and on other networks where the field remains unpopulated, would unblock the live-crowding feature this project could not build.
- **Audit-and-republish is the realistic way to keep content current today.** Sensory content decays as venues change. Aspect's twice-yearly Quiet Room re-audit cadence (§9.1.2) is the available precedent for on-cycle revisions, with out-of-cycle revisions triggered by operator-flagged venue changes. The cycle must be cheap enough to run on the cadence venue change demands. A baseline cadence is set out in §10.

9. Framework & Design Guidelines

This section sets out the Sensory-Sensitive Wayfinding Framework and the design guidelines evidenced by the focus groups and the Adelaide Railway Station build: six framework components, 22 design guidelines that operationalise them, and three open questions.

The evidence base comprises two Milestone 3 focus groups (six lived-experience participants across FG1 in March 2026 and FG2 in May 2026, three Aspect auditors in FG2 admin testing), the Milestone 1 problem-research interviews, the Milestone 2 sensory-map validation focus group, and one piloted venue.

Transferability claims pertain to other Autistic-cohort venues and similar accessibility-content contexts. The constituencies left untested at §3.5 (dementia, low vision, PTSD, generalised anxiety, hearing impairment, non-autism sensory processing differences) are carried as a research-extension recommendation in §10.

9.1 The Sensory-Sensitive Wayfinding Framework

The framework has six components. Each was named in the focus-group evidence and traced to design or method choices in the build. Synthesis into six was an interpretive task done after the sessions.

9.1.1 Place, subject, detail hierarchy

Content sits in a three-level tree: venue (place), area or facility (subject), per-category sensory detail. This matches the planning task: a visitor chooses a venue, locates their area, then sees only the relevant detail.

Both sessions named the within-page hierarchy:

> *I did struggle to understand all of the information […] my pattern recognition was kind of jumping between things and not quite sure on like the hierarchy of what to start with first and follow to. (FG1-P1, Focus Group 1)*
>
> *having this information higher up in the hierarchy of information potentially as well, right? Like that could be a better way to structure it just so that like you got your key stuff at the top like. Quiet zones, bathrooms, entries, exits, whatever. (FG2-P1, Focus Group 2)*

Prior-art review (§2.4, Appendix D) found successful venues now publish below venue level: NGV per exhibition, MCG per event type, Stratford Festival and Arts Centre Melbourne per production, Canberra Airport per direction.

The tree accommodates sub-resource publishing. The Adelaide Railway Station build does not exercise this at scale, and whether the admin tool can drive it remains untested.

Design response. The Sensory Guide reads as place, subject, detail: one public landing page per venue, ordered area sections, per-category sensory details inside each area (§4.3.1). Key facilities sit at top and bottom of the page above area-level prose.

Transfer. Venues with a single relevant space collapse the tree to two levels. The pattern is expected to transfer, with grain set venue by venue.

9.1.2 Before, during, after journey phasing

Resources should support two user phases (planning and in-venue use) plus a post-visit operational loop closing the content-currency cycle. Planning was dominant in this cohort, in-venue use secondary.

Planning came up unprompted as the primary use case in both sessions:

> *It reduces my anxiety to be able to picture and anticipate what's going to happen […] it takes away a lot of the anticipatory anxiety. (FG1-P1, Focus Group 1)*

> *it reduces my anxiety and makes me feel better about being out in public because it reduces the layer of unpredictability. (FG2-P1, Focus Group 2)*

In-venue use is secondary, not absent:

> *Yes, but more likely when I arrive unless the venue has told me this exists. (FG2-P2, Focus Group 2 chat)*

Focus Group 2 admin testing found the feedback affordance correctly placed but under-prominent.

Aspect's Quiet Room accreditation criteria commit the organisation to collecting, acting on and sharing feedback (Aspect, accessed 19 May 2026), and accredited rooms are re-audited twice yearly. That cadence is the framework's reference point for re-audit, and supports a feedback path on the admin side without elevating it to a third user phase.

Design response. The same authored content serves both phases plus the operational loop: stable public URL, responsive mobile layout, in-page feedback control routing into the version-history and last-updated controls (§4.3.1, §4.4.14).

Major venue changes trigger an out-of-cycle update. The framework recommends naming the cadence at venue level.

Transfer. The two-phase frame is expected to hold across venues. Phase weighting is cohort-dependent, and this project tested only the planning-dominant cohort.

9.1.3 Category as filter axis, with severity as a secondary configurable layer

Two filter dimensions sit at different maturities.

The *category* dimension tested cleanly across both focus groups. The focus-group-tested core is sound, light, smell, crowds and temperature. The platform supports four additional categories (touch, movement, vibration, air quality) installed at Adelaide Railway Station (§6.2). The dimension is settled.

The *severity* dimension (Low / Moderate / High or venue-overridable equivalents) is in flight. No Australian standard exists, Aspect's terminology was pending at time of writing, and Focus Group 2 flagged the current labels as needing a clearer rubric.

A reader can scope to one category, several, or 'high severity only'. Melbourne Airport publishes three-tier sensory zoning across its four terminals.

The category dimension tested well:

> *I'm not great with sound and light, crowds I can deal with and smell I can't deal with. So I like the categories. I think that's quite a good representation of it. (FG2-P1, Focus Group 2)*

> *I love that you could filter by the senses. That was really helpful. So I just put the key senses that I'm impacted by and went through by those. (FG1-P1, Focus Group 1)*

Severity wording remains unresolved:

> *If they have a rating they used to decide I want to know it. But if it's vibes that's fine. (FG2-P2, Focus Group 2 chat)*

Design response. Multi-select category filters auto-expand matching sections (§4.4.1). Per-area severity badges and per-detail level use venue-overridable labels (§4.4.7).

Transfer. Categories are expected to transfer cleanly. Severity wording is configurable because no Australian standard exists.

FG1-P3 flagged that smell varies more across time than other inputs, so smell entries should be authored conservatively or marked as time-variable. Unresolved at §9.3.

9.1.4 Multi-format delivery as access requirement

Multi-format output is not a packaging convenience. Users prefer different artefacts, and institutions only accept certain ones.

The user-side case:

> *often, yeah, we do need to provide information in different ways or gather information in different ways. And so being able to have that flexibility and options is really valuable. (FG2-P3, Focus Group 2)*

The institutional need came from PTSA correspondence (§6.4): a Word document for the customer-facing template, plus a PDF for the accessibility and inclusion resources suite. Procurement and security policies drove the format requirement.

Prior-art evidence reinforces the move. Static PDF dominates the corpus. Web-first formats (Australian Museum, Tate Modern, MCA, Monash Libraries, DFO Brisbane) are exceptions, and audio is rarer still (V&A's audio descriptions and tactile-with-QR Raphael Cartoons).

Parks Victoria's editable DOCX plus Key Word Sign Australia versions are the closest multi-format precedent. Single-PDF as the corpus default fails too many users and too many partners.

Design response. Every published guide produces five artefacts from one source: an interactive web page at a public URL, a web component for placement inside a partner's own site, a filterable PDF with QR codes substituting for the interactive embeds and maps, a Word document, and an MP3 audio narration cached per published version (§4.4.3).

Transfer. The five-artefact set is the starting point for government, transport, and large-institution partners. Smaller venues may need a subset. The one-source-many-outputs principle is expected to transfer; the artefact list is venue-context dependent.

9.1.5 AI as suggester, not author

AI surfaces suggestions (flagging subjective language, prompting for missing photos, summarising) but does not silently rewrite authored content.

Focus Group 1 named the inversion:

> *I think sometimes it can be helpful to check our work rather than changing our work and us checking it. (FG1-P5, Focus Group 1)*
>
> *I think if it we were cautious of it not changing language, but even just identifying, hey, we've noticed an error, would you like it to change it to this might be really good. (FG1-P1, Focus Group 1)*

Focus Group 2 testing the rebuilt admin tool surfaced no objection to the inversion. Absence of objection is the available signal, short of direct endorsement (§4.4.11):

> *It's sort of like spell check. I would go, oh, actually that is, I did spell that wrong. […] I will accept that. (FG2-P3, Focus Group 2, accepting a suggestion that flagged 'hot' as subjective)*

Two trust scaffolds keep the inversion honest. Every suggestion has accept and dismiss controls and shows the proposed replacement as a strike-through diff (§4.4.11). The PDF-import path produces a document-level checklist above the draft preview (§4.4.16).

Trust also depends on disclosure. FG1-P1 asked in chat which AI tool the system uses and flagged that internet-trained models can carry problematic language.

FG1-P3 said: 'I know some people in our team use AI for different reasons. Personally, I try to avoid it wherever I can. It makes me feel very uneasy for different ethical reasons.'

The model should be disclosed. Training-data concerns are discussed in §10.

Design response. AI in the Sensory Guide proposes changes that the author accepts or rejects (§4.4.11).

Transfer. The pattern fits accessibility-content authoring, where content authority must stay with a human expert and the cost of an AI-introduced mistake exceeds the cost of human review of every change. Generalisation depends on that cost geometry. Other patterns (AI drafts with audited sign-off) suit domains where mistakes are cheaper to catch downstream.

9.1.6 Lived-experience-led research method

Autistic participants shaped the design while it was being built. Their feedback drove the decisions, instead of being gathered at the end to sign off work that was already finished.

The framework was developed through iterative focus-group research with Autistic lived-experience participants. Aspect's Autism Friendly team recruited, co-designed the question sets, ran the sessions, and provided partner review between iterations.

Between the two Milestone 3 sessions the design changed direction (map-to-guide pivot at §3.4 and §5.4): the map-led prototype was replaced with the guide-led Sensory Guide, and AI moved from author to suggester. Both pivots came from lived-experience feedback.

> *the information needs to be developed based on an assessment of the venue by people, by autistic people. (FG1-P5, Focus Group 1)*

Aspect's Autism Friendly team includes Autistic staff who bring lived-experience perspective alongside professional expertise.

Design response. Every product decision in the §4 catalogue traces to a session quote and a §5 iteration change.

The pilot tested an authoring constellation of Aspect auditor plus venue content-supplier plus AI suggester. Whether the framework works under different constellations, and whether the same six components surface from research with other sensory-sensitive constituencies (Deafblind, dementia-affected, PTSD, generalised anxiety, hearing impairment, non-autism sensory processing differences and older adults), is untested and carried in §10.

Transfer. On this project's evidence, the people with lived experience need to help make the design decisions, rather than only checking them once the work is done. Three operating choices made it practical:

- partner with an organisation that brings lived experience and access to the constituency (Aspect here),
- compensate participants for their time, in line with sector practice,
- review between iterations rather than at the end.

9.2 Design guidelines

22 design guidelines are distilled from the focus-group evidence. Each traces to a verbatim participant quote and to either a built Sensory Guide feature or a flagged §10 item.

Guidelines are grouped public-side first, then content, then admin-side. Each entry follows a three-part structure: principle statement, evidence quote, and *Built as* / *Transfer* lines.

Public-side guidelines

Hierarchy and disclosure

1. Photographs over schematic maps. Real photographs sit higher in the page hierarchy than the map.

> *my brain is always trying to like foresee and plan everything in the future. So if I've got more prep and predict information like clear photos, they really, really helped me. (FG1-P1, Focus Group 1)*

Why photos rather than maps in the lead position: see the map-as-overwhelming evidence at G8.

Built as: per-area photo carousel and full-screen lightbox (§4.4.6). *Transfer:* venues with photograph-able interior spaces, rarer for outdoor or changing environments.

2. Progressive disclosure on cumulative sensory load. Sections start collapsed. Visitors expand only what they want.

> *I thought the information was great, it was just kind of too much at once and then I was overwhelmed and then I couldn't really like piece it through. (FG1-P1, Focus Group 1)*
>
> *maybe that needs to be not expanded on initial load […] padding it out just means that you're going to be scrolling and it just feels unnecessary. (FG2-P1, Focus Group 2)*

Built as: collapsed-by-default area sections, expand-all / collapse-all controls, and a collapsible 'how to use' panel (§4.4.2). *Transfer:* content-dense interfaces.

3. Sensory information is opt-in, not the default first view. Different users want different amounts, and a fixed depth fails both ends (§3.4 contradictory-preferences finding).

> *I do think the list of how to use is useful, but that's something I'd put under a drop down (like [FG2-P1] said). (FG2-P2, Focus Group 2 chat)*

Built as: collapsed defaults, multi-select filters, persistent filter state (§4.4.1, §4.4.2). *Transfer:* content-dense interfaces serving heterogeneous readers.

4. Key facilities sit high in the page hierarchy. Quiet zones, accessible bathrooms, entries, and exits are the most-used items and should not be buried in area-level prose.

> *put everything else that someone might need (eg maps and where quiet spaces are) earlier so they don't bounce off whole website. (FG2-P2, Focus Group 2 chat)*

Built as: auto-generated about-this-guide overview and Key Facilities block (§4.4.4). *Transfer:* content-dense venue resources.

5. Quiet-space discovery is a first-class concept. Quiet spaces were the primary need FG2-P2 expressed, and prior-art (Melbourne Airport's three-tier zoning, Aspect's Quiet Room scheme) confirms the pattern at scale. They should appear as a top-level filter or facility, above a generic 'key facilities' list.

> *The only thing I would want from the sensory information is where the quiet areas are. (FG2-P2, Focus Group 2)*
>
> *I think the quiet spaces one could be another button to identify. (FG2-P3, Focus Group 2)*

Built as: partially built. Quiet zones appear in the Key Facilities block, with top-level filter promotion flagged for §10. *Transfer:* wherever a venue has designated low-stimulus spaces.

Filters and interaction

6. Filtered sections auto-expand. When a reader applies a category filter, matching sections expand automatically. Collapse-by-default holds for unfiltered content. Tested in Focus Group 1, accepted in the rebuilt Focus Group 2 behaviour.

Built as: category-filter auto-expand (§4.4.1). *Transfer:* filterable content-dense interfaces.

7. Filter UI uses familiar conventions; bespoke interaction pays a discoverability tax. Google Maps is the mental model this cohort already knows.

> *Google Maps is just the software I tend to use, so I naturally kind of went to use the software as if it was Google Maps and I did get lost in it and get confused at different points. (FG1-P1, Focus Group 1)*
>
> *designing as close to the familiar as possible would be helpful. (FG1-P6, Focus Group 1)*

Built as: the redesign removed the separate map mode and added an in-line 'Show on map' button per area (§4.4.5). *Transfer:* wherever a familiar pattern exists for the task.

Visual language

8. Icons should not sit on a navigational map for sensory-sensitive users. The map-led prototype (icons overlaid on the navigational map) drew the strongest 'overwhelming' feedback in this cohort.

> *a sensory map currently could be extra confusing because there's a lot of visual information. Some of it is a new language to learn. (FG1-P5, Focus Group 1)*

> *I did feel pretty overwhelmed looking at that map. I felt like it was a new level of, um, complication. (FG1-P6, Focus Group 1)*

Built as: the redesign removed icons from the map and represented sensory categories as text badges next to area headings (§4.4.7). *Transfer:* any navigational map where symbolic overlays compete with wayfinding.

9. Visual category colour-coding is consistent across map, badges, and prose. Same category, same colour, everywhere.

> *maybe color coding the categories because they're color coded as you can see… and then they're not color coded here. I just feel like that would be an easy way for my brain to go, oh, that means that, that means that. (FG2-P1, Focus Group 2)*
>
> *I like the consistent colours, that the sensory details have their own colour. (FG2-A1, Focus Group 2 chat)*

Built as: design-system tokens used by category badges, the sensory key, and per-detail labelling (§4.4.7). *Transfer:* multi-view systems using categorical encoding.

Cross-cutting

10. Where video is offered, prefer 360-degree walkthroughs over linear video.

> *The VR walkthrough style thing is best because you can do 360 degrees. You can walk like plot your route through instead of having to click on the dot, go in, come back out to the map, click on the next dot. (FG1-P5, Focus Group 1)*

FG1-P3 preferred the Google Street View model: 'swivel the camera around so I can see what the buildings next door look like.'

Built as: not yet built, flagged as a §10 post-MVP item. *Transfer:* wherever capture is feasible.

11. Last-updated date is prominent; a changelog should accompany it. Trust depends on content currency.

> *I liked the last updated thing at the top. I think people knowing when it was updated is really helpful and important. (FG1-P5, Focus Group 1)*

In Focus Group 2, FG2-P3 asked for a human-readable changelog per version.

Built as: last-updated and version history built, per-version diff flagged for §10. *Transfer:* wherever venue content changes between audits.

12. Mobile-first works for both planning and in-venue checks. A common design assumption maps planning to desktop and in-venue use to mobile. Focus Group 2 evidence contradicts it. Participants used mobile across both phases, including for pre-trip planning.

> *I would probably have it on my phone is the first thing […] is this web safe for a mobile phone environment or whether I'd have to download the app. (FG2-P1, Focus Group 2)*
>
> *I agree, I do it in advance, but I would do on mobile too. (FG2-P2, Focus Group 2 chat)*

Built as: one responsive codebase from phone through desktop (§4.4.8). *Transfer:* device-to-phase mapping is a useful lens rather than a fixed rule.

Content guidelines

13. Content uses objective, stimulus-focused language, not subjective. Describe what is present, not how a user is expected to feel about it.

> *There are a few examples of quite subjective language that I think could be just made more objective and simpler. So things like noticeable temperature, comfortable, uncomfortable, unbearable smell of ammonia. (FG1-P3, Focus Group 1)*
>
> *I really liked the simple but descriptive language when it didn't say how the person felt […] 'The sound of trains travelling overhead creates a loud, deep rumble that reverberates strongly through the metal tunnel.' […] I could picture the sound. (FG1-P1, Focus Group 1)*

Built as: admin tool's 'remove subjective wording' AI-suggestion behaviour against a banned-word list (§4.4.11). *Transfer:* wherever stimulus-based content is authored.

14. Lighting descriptions are specific (warm vs cool, not generic). Stimulus-specific descriptions help users predict their response.

> *when we speak about lighting, I would've personally liked to say whether it was warm or cool lighting. (FG1-P1, Focus Group 1)*

Built as: not yet enforced by the suggester, flagged for the simplify and subjective-wording prompts. *Transfer:* wherever stimulus categories admit a sub-distinction.

15. All images carry alt text. Standard accessibility practice.

> *I don't know how this works, but it would be great if the images can have alt text. (FG2-A1, Focus Group 2 chat)*

Built as: added to the photo-management tools, with verification at publish time next. *Transfer:* WCAG baseline.

16. Severity wording is configurable, not standardised. Configurable labels accommodate partner-specific terminology while no shared rubric exists.

Built as: venue-overridable severity labels (§4.4.7; underlying finding at §9.1.3). *Transfer:* default until a standard exists. The recommendation at §10 is that Standards Australia adopt a sensory-severity terminology standard, with related work on colour scales and icon sets.

17. Identity-first language matches the constituency convention. Where the partner constituency is the Autistic community, use 'Autistic people' (capital A, identity-first), mirroring Aspect's materials.

> *a world where Autistic people, their families, and our staff can grow, connect, and thrive. (Aspect vision statement, verbatim)*

Built as: adopted in the design-system glossary. *Transfer:* wherever the constituency has a stated language preference. Check with the partner organisation.

18. Terminology distinguishes access modes ('mobility access route', not 'accessible route'). 'Accessible' is overloaded across disability types and obscures which mode is meant.

> *the naming convention for accessible routes […] I just think maybe in the future accessible routes won't be the most appropriate term because accessible means different things for different disabilities. So I would probably say like mobility access route. (FG1-P5, Focus Group 1)*

Built as: adopted in the design-system glossary. *Transfer:* wherever route-naming distinguishes modes.

Admin-side guidelines

19. AI suggestions are acceptable or rejectable; never automatic edits. The user must consciously accept each suggestion. The evidence is at §9.1.5 *AI as suggester, not author*.

Built as: suggestion card with accept and dismiss controls (§4.4.11). *Transfer:* authoring workflow where content authority must stay with a human expert.

20. Publish-readiness warnings are shown to admins on missing content. Active prompts raise quality without demanding admin diligence.

> *there's a bit of a warning here. There's some things missing, so there's no photos. I love that as well because it reminds me of what I need to still do. (FG2-P3, Focus Group 2)*

Built as: publish-readiness blocks (publish disabled until every selected category has both a level and description) and the post-import suggestion list (§4.4.16). *Transfer:* authoring workflows with a completeness contract.

21. Voice and dictation input on long-form authoring fields. Reduces friction for admins capturing content on-site, especially Autistic admins who may prefer voice notes.

> *Actually that would that would be great if you could speak into the app and then it could transcribe it and then automatically put it in. (FG1-P1, Focus Group 1)*
>
> *I like that you can dictate during that section because I know a lot of us like to make voice notes on our phone. (FG2-P4, Focus Group 2)*

Built as: long-form fields with cloud transcription and a live input-level indicator (§4.4.9). *Transfer:* authoring workflows with on-site capture.

22. AI model is disclosed to users who ask. Trust requires disclosure (see the disclosure evidence at §9.1.5).

Built as: model-identity disclosure is in scope for §10 product-documentation work, with the related training-data concern carried separately in §10. *Transfer:* any AI-assisted product whose users have a stake in the content.

9.3 Open framework questions

Three open design questions remain. Each is a candidate for further research.

- **Map-to-section linkage in the reverse direction.** The per-area ‘Show on map’ control (section to map) was Focus Group 2-validated. The inverse (map marker to section) was not tested and remains open.
- **Photograph framing.** Focus Group 1 split on framing. FG1-P1 wanted zoomed-out context for orientation, FG1-P5 wanted first-person entry-direction framing. The current build leaves framing to the auditor, and whether a guideline can be drawn that serves both is unresolved.
- **Smell as a filterable category, accuracy caveat.** FG1-P3 flagged that smell varies more across time than the other sensory inputs, raising a question of how to keep filter-by-smell information current. The filter remains category-correct, but the underlying entry should be marked as time-variable. Practice convention for time-variable entries is unresolved.

10. Recommendations for Broader Deployment

10.1 The validated set

A subset of this project’s design choices is well-supported by the focus-group evidence and the pilot at Adelaide Railway Station.

- **Sensory-map to planning-guide pivot.** Map-led interfaces tested as a source of overwhelm in Milestone 2 and Focus Group 1; the guide-led redesign tested in Focus Group 2 was received positively, with overwhelm language absent (§3, §5.2, §5.4.4, §9.1.1). Focus Group 2 did not re-test the original artefact, so the comparison is reception-of-the-redesign rather than a controlled before-and-after (§7.3).
- **Place-subject-detail hierarchy with sub-resource granularity.** Venue → area → per-category detail is the right tree for this cohort's planning task and matches the prior-art pattern of below-venue publishing (per exhibition, per event type, per direction) (§9.1.1, Appendix D).
- **Progressive disclosure as the default for cumulative sensory load.** Collapsed-by-default sections with expand-on-tap, expand-all / collapse-all, and auto-expand on filter selection addressed the overwhelm Focus Group 1 reported (§4.4.2, §5.4.2).
- **Category filters as the primary way visitors personalise the guide.** Multi-select category filters with persistent state served the dominant 'choice and control' theme and resolve the information-density contradictions in §3.4 (§4.4.1, §9.1.3).
- **Multi-format publishing from one source.** Web, embeddable partner version, PDF with QR codes, Word, and audio from a single draft. Government and enterprise procurement made multi-format mandatory. The PTSA Word-document request is the worked partner instance (§4.4.3, §6.2, §9.1.4).
- **AI as suggester, with human authorship preserved.** Suggestion cards with accept and dismiss controls, strike-through diffs, and short rationales. Tested positively in Focus Group 2 against the rebuilt admin tool (§4.4.11, §5.4.1, §9.1.5).
- **Real photographs at higher priority than the navigational map.** Per-area photo carousels with captions and a lightbox. Photographs reduced anticipatory anxiety across both focus groups (§4.4.6, §9.2 G1).
- **Familiar UI conventions, no bespoke map interaction.** The separate map mode was replaced by an inline per-area 'Show on map' control. The rotation-recovery issue raised in Focus Group 1 disappeared with the redesign (§4.4.5, §5.4.4, §9.2 G7).
- **Mobile-first authoring with voice dictation and on-device capture.** On-site mobile capture, voice on every long-form field, and floating-action camera and microphone buttons matched the workflow FG1-P5 described in Focus Group 1 (§4.4.9, §4.4.10, §5.4.3).
- **Operator-driven content-currency loop with version history.** Per-venue version history, revert, and rename-flags-unpublished support the operator-driven re-publish workflow end-to-end (§4.4.14).
- **Automated accessibility, performance, and security checks on every code change.** Pa11y, Lighthouse, and dependency-audit checks run in the deployment pipeline and block release on failure (§7.5).

- **Aspect-led lived-experience research as the design input layer.** Recruitment, co-designed question sets, sessions, and review between iterations all routed through Aspect's Autism Friendly team (§3, §9.1.6).

10.2 Next-stage research

Three research questions sit outside the Milestone 3 evidence base and are the priorities for the next stage of work. Severity-wording standardisation is taken up in §10.3.

- **A pre-and-post planning-anxiety and confidence instrument.** Direction of effect is consistent across the participant accounts (§7.3); magnitude is not. A short repeatable measure administered before and after a piloted deployment would quantify what both focus groups reported qualitatively. The instrument should sit alongside in-product analytics.
- **Transfer to other venue types.** The pilot tested one underground rail station. The framework should be tested at airports, shopping centres, stadiums, universities, civic spaces, museums, and hospitals. Existing prior-art coverage varies (§2.4, Appendix D).

  Healthcare is the standout prior-art gap and the only candidate carrying a named federal policy commitment (the National Autism Strategy 2025–2031's sensory-friendly health undertaking). No Australian hospital publishing a public sensory map, social story or action plan could be located during the prior-art review described in Appendix D. Hospital-passport references appear in NSW Health and Sydney Children's Hospitals Network policy material but the artefacts themselves were not publicly accessible. The other venue types sit alongside healthcare as candidates without the same policy hook.
- **Transfer to other constituencies.** Recruitment through Aspect concentrated coverage on Autistic lived experience by design. Transfer to other sensory-sensitive constituencies (Deafblind, dementia-affected, PTSD, generalised anxiety, hearing impairment, non-autism sensory processing differences, older adults) is untested.

  Each would need its own partner organisation (Dementia Australia for the dementia cohort, Vision Australia for low vision) and a recruitment route comparable to Aspect's.

The framework is valid for this cohort and venue. The per-component transferability claims are stated in §9 as project hypotheses, and need further validation.

10.3 Australian standards gap

No Australian standard exists for the visual language sensory-mapping resources use.

- **Severity wording and definitions.** FG1-P5 flagged the default Low / Medium / High labels as too thin and partly subjective, and pointed to Aspect's internal definitions as a starting point. Focus Group 2 repeated the concern (§5.4.5). Aspect's preferred wording remained pending at time of writing.

- **Colour scales.** Severity colour-coding was tuned against WCAG 2.2 AA contrast and adjacent-level distinguishability inside this build (§5.4.4). No national reference exists to align it to.
- **Icon sets.** The sensory-map prototype's icon-on-map overlay generated the strongest 'overwhelming' feedback in Focus Group 1 (§3, §5.4.4). Removing icons from the navigational map addressed the symptom in this build. A shared icon vocabulary for sensory categories does not exist nationally.

The recommendation is that Standards Australia, Aspect, Autism SA's Autism Friendly Charter and Dementia Australia collaborate on drafting an Australian conformity standard for sensory mapping. Aspect's Autism Friendly Accreditation applies specifically to certified Quiet Rooms, not to venues at large. The standard would cover severity terminology, colour scales, and icon sets.

The three certification frameworks operating in Australia (§2.4) are not cross-recognised. A shared visual-language standard would let partners author against one vocabulary regardless of which certification they pursue.

The case for the standard is well-supported by this project's evidence and the prior-art corpus reviewed in §2.4 and Appendix D. The standard's content is a separate piece of work.

10.4 Platform and partner deployment

Cross-venue personalisation

A profile of category and severity preferences stored on the user's device and applied across every venue they visit. Filters persist per venue today. Cross-venue persistence is the next-step personalisation, well-supported by the 'choice and control' theme.

Partner network cascade

The pilot's partnership configuration (Aspect for audits, PTSA for operator content, BindiMaps for platform) can extend through the partners' existing networks.

- **Aspect commercial working relationship.** At the end of Focus Group 2 an Aspect facilitator proposed a formal sensory-assessment service, with Aspect running audits and BindiMaps providing the publishing platform. The BindiMaps team signalled reciprocal interest.

  Aspect's Autism Friendly team consults with organisations across transport, retail, cultural institutions, councils and sporting venues (§3.2). The recommendation is to formalise that proposed working relationship. The direction is well-supported. Commercial terms are out of scope here.
- **PTSA operator network.** Adelaide Metro's wider station footprint and the cross-state transit-operator network are the natural next venues after Adelaide Railway Station. PTSA's procurement experience inside this pilot (including the Word-document and PDF

distribution paths) is reusable for comparable operators, and is well-supported by the pilot evidence. Uptake by other operators is the next test.

- **Venue-subscription operating model.** A venue-subscription arrangement (per-venue editor licensing, with audit services procured separately through Aspect or comparable accredited assessors) is one operating shape evidenced by the PTSA engagement. This is a forward signal rather than a delivered outcome.

Pre-deployment audits

Two audits should be commissioned before the admin tool is opened to external editors beyond the project's known set.

- **Independent third-party WCAG sign-off with assistive-technology session testing.** The 19 May 2026 accessibility audit cleared the build for release with minor findings only, but §7.5 records the caveat that the audit was automated and is not independent.

  Real-user sessions on screen readers (NVDA, JAWS, VoiceOver, TalkBack), switch input and voice control, alongside an export-artefact audit (PDF, DOCX, MP3), would carry the audit out of automated scope.

- **Formal independent penetration test plus an AI prompt-injection threat model.** The 19 May 2026 security audit cleared the build for continued deployment subject to remediating two high-severity findings, both fixable in under 30 lines of code (§7.5).

  An independent test would carry the audit out of the automated-tooling caveat. It would also let the AI-driven features (PDF transform, suggestions, overview, audio) be assessed for prompt-injection specifically.

Discovery and promotion

Discovery is critical. FG2-P3 raised the discoverability question explicitly at the end of Focus Group 2, asking how visitors would find the guide in the first place. Promotion routes identified by the pilot:

- the venue's own customer-information channels (PTSA's Customer Strategy and Inclusion Lead flagged the Adelaide Metro website as the longer-term distribution path, §6.4);
- Aspect's organisational network (§3.2);
- the support-worker and carer-worker constituencies FG2-P1 and FG2-P3 named as a use case alongside end users.

Discovery is well-supported as a deployment requirement. Channel mix sits with the operating team.

10.5 Data-availability recommendations

The data-dependency findings in §8 carry three recommendations for venue and transit data publishing.

- **Australian transit operators should publish `OccupancyStatus` in GTFS-RT.** The field is structurally available in the protocol for any operator that chooses to populate it. Transport for NSW already publishes it. Adelaide Metro does not (§8.2).

  Closing the gap for the Adelaide network and other Australian networks would let the live-crowding feature this project could not build proceed. The recommendation is addressed to Adelaide Metro and to comparable operators, with Transport for NSW's existing implementation as the working reference.
- **Venue-side IoT integration is a partner question, not a platform question.** Where a venue independently deploys sound-level, air-quality, or other sensors, the admin tool can ingest the readings. The blocker is hardware presence at the venue. Venues that deploy IoT for operational reasons (heating-ventilation, security) can extend the same deployment to feed sensory-relevant signals into the guide.
- **Maintenance and cleaning schedules can flow through per-venue data-sharing arrangements.** Internal operational systems have no public APIs (§8.2). A per-venue integration with a willing operator is feasible and would let the maintenance-and-cleaning feature this project could not build proceed for that venue. It is not a framework-level capability until a sector-wide data-sharing pattern exists.

Confidence on all three: well-supported by this project's feasibility findings (§8). The action lies with operators and venues rather than with the product.

10.6 Post-MVP feature recommendations

A focused set of feature recommendations from the iteration evidence pack, grouped by where they apply (public-side, content, admin-side). Maturity is recorded as **researched** (specified and partly designed but not built), **prototyped** (partially built but not released), or **not started**, anchored to the roadmap horizons (researched ≈ 12-month, prototyped ≈ 24-month, not started ≈ 36-month). The list contains only items with a named focus-group source; further ideation items without named participant sources are excluded.

Public-side

Feature	Source	User need	Maturity
Embedded cafe and restaurant menus	FG2-P1; FG2-P3 followed up; FG2-P2 agreed	Pre-visit planning to reduce in-moment decision load	Researched (platform supports generic embedded media; menu-domain

Feature	Source	User need	Maturity
			not yet)
Vending-machine, water, and coffee discoverability as a top-level category	FG2-P1 on water; FG2-P4 added coffee	Self-regulation supplies for sensory-load management	Not started as a dedicated category
Walk times between areas	FG1-P6, post-FG1	Journey planning, first-time visitors	Not started
Peak and busy times for busy areas (static, historical)	FG1-P1 (post-session); FG2-P2 asked it be raised higher	Lets users actively choose less-busy times; live crowding infeasible (§8), so static historical advice authored by the auditor	Prototyped (per-area busy-times block exists, §4 P15); peak-times authoring guidance and prominence elevation not started
Quiet-space prominence as a top-level facility or filter	FG2-P2 (most consistent thread): ‘The only thing I would want from the sensory information is where the quiet areas are’; FG2-P3 agreed	Cross-references prior-art evidence at scale (Melbourne Airport’s three-tier sensory zoning, Aspect’s Quiet Room scheme); currently inside Key Facilities (§4.4.4), recommendation is to elevate	Prototyped at lower prominence
Update history and per-version changelog	FG2-P3	Human-readable changelog per published version supports trust in content currency	Prototyped admin-side; not started public-side
Hyperlinking inside the public guide	FG2-P1: tap a facility in the summary, jump to the matching area section	Navigation inside long pages for the ‘less info, key things only’ cohort	Not started

Feature	Source	User need	Maturity
Category colour-coding inside body prose	FG2-P1: where prose names a category, colour it in the matching category badge colour	Visual scanning for the ‘skim, don’t read’ cohort	Not started
Default-collapse the how-to-use section, with collapsible sub-headings	FG2-P1, FG2-P2, and FG2-P3 in FG2; unanimous in the room	Returning-user load reduction	Not started
Collapsible per-category sub-sections inside an opened area	FG2-P1 and FG2-P2	Progressive disclosure for the ‘less info’ cohort while preserving access for the ‘all info’ cohort; the within-section granularity FG2-P2 pointed to in §4.4.2	Not started
360-degree and virtual-walkthrough photos	Universal endorsement (FG2)	Familiarisation pre-visit, especially for the photo-loving cohort	Prototyped at the platform layer (BindiMaps); guide-UI integration not yet built

Content

Feature	Source	User need	Maturity
Editable AI-generated overview and key-facilities text	FG2-P3: ‘What happens if it generates something that isn’t accurate, for instance, that would need to be edited’	Auditor control over accuracy, a ‘choice and control’ theme item	Prototyped (regenerate, §4.4.4); not started (edit-in-place)

Feature	Source	User need	Maturity
Per-image alt text in the admin	FG2-A1	Screen-reader accessibility (WCAG baseline)	Not started
Embedded-media heading and descriptor	FG2-P3: a generic ‘Show embed’ button gives no context	Comprehension and accessibility before opening	Not started
Lighting descriptions sub-classified (warm versus cool)	FG1-P1	Addition to the simplify and subjective-wording prompts	Not started in the suggester

Admin-side

Feature	Source	User need	Maturity
Email-automated multi-admin invite	FG2-A1	Smoother multi-admin collaboration; MVP step generates a link plus copy-able message text (§4.4.12)	Prototyped at the link-plus-script step
Editor-feedback notification emails	Implicit from FG2 notes; participants did not find the feedback button until late in the session (§4.4.13, §7.4)	Continuous-improvement loop closing	Prototyped on submission; not started on notification
Bulk and drag-and-drop photo upload	FG2-P3 and FG2-A1 in FG2; drag-and-drop failed in-session	Different audit workflows want different capture orders	Prototyped (per-area upload works); drag-and-drop is partial
Reorder area-editor tabs so the sensory tab precedes the summary tab	FG2-P3 and FG2-A1 (FG2)	Match the natural on-site auditing order	Not started

Feature	Source	User need	Maturity
Enter-to-submit on the admin map URL field	FG2-P3	Standard form UX	Not started
Visible next-buttons in the area editor (above the fold on mobile)	FG2-P3: did not see them until halfway through the session	Discoverability for first-time admins	Prototyped (buttons exist); layout fix needed
Reverse map-to-section linkage	Built in inverse direction (section → map); inverse untested	The per-area ‘Show on map’ control is built and FG2-validated; map-marker → section remains an open framework question (§9.3) ahead of build	Not started (open question)

Confidence on the feature list: maturity statuses are accurate as of 19 May 2026.

10.7 Residual risks and continuity

Four risks remain at handover. Three are addressed inside §8 and earlier in §10; the fourth carries across all of them.

- **Data dependency.** Four dynamic-content feeds are unavailable for Adelaide Metro at the time of writing: live transit-crowding, sound-level IoT, maintenance and cleaning schedules, and Google Popular Times (§8).

  The MVP carries the dynamic-content need through operator-supplied content. The audit-and-republish cadence is the realistic way to keep content current today (§8.5, §9.1.2). The §10.5 publishing recommendations would close the gap. An operating team should monitor `OccupancyStatus` availability across Australian operators.
- **Partner adoption.** PTSA and Aspect are the demonstrated partners for the pilot. Broader uptake across each network is a deployment question (§10.4) rather than a research one. A partner adoption that loses momentum carries the discovery and promotion risk.
- **Content governance.** Sensory content has a venue half-life. Out-of-cycle revisions are triggered by operator-flagged venue changes.

Aspect's twice-yearly Quiet Room re-audit cadence (§9.1.2) is the available precedent for on-cycle revisions. The cycle's running cost falls on the operating team and the auditing partner. The platform's role is to keep it cheap enough to run.

- **Post-pilot continuity.** The operating team that carries the product forward inherits the three risks above and the pre-deployment audits at §10.4. The technical-debt position at handover is documented in §7.5 and Appendix B.

11. Conclusion

11.1 Findings

For this cohort at Adelaide Railway Station, pre-trip planning came up unprompted as the primary use case in both focus groups. The organising finding behind it was choice and control: unpredictability matters as much as the stimulus, and the anticipation and choice a person has determine whether the same environment is navigable or overwhelming.

Lived-experience review between iterations shaped the build as it went, catching two design pivots before launch. Those findings, with others, reshaped the product:

- **Cumulative sensory load on map-led interfaces.** The map-led sensory interface tested as a source of overwhelm in Milestone 2 and again in Focus Group 1. The redesigned Sensory Guide tested in Focus Group 2 was received positively, with overwhelm language absent in discussion of the redesign.

  Progressive disclosure, per-category filters and per-area 'Show on map' buttons replaced map-led density.
- **Photographs over schematic maps.** Photographs ranked higher than schematic maps as the mental model participants used to picture themselves in a venue. The guide ranks photos above the venue map in the page hierarchy.
- **AI as helper, not author.** Silent AI authorship in the admin tool was rejected in Focus Group 1. The rebuild has AI propose sentence-level edits with a short rationale that the auditor accepts or dismisses.

Contradictory preferences across participants emerged as a finding in their own right. Examples from the focus-group data:

- FG2-P1 and FG2-P2 disagreed on information density on the same overview page.
- Focus Group 1 split between map-anchored and text-anchored entry points.
- Planning-heavy and arrive-and-look-it-up patterns coexisted in the cohort.

Where participant needs span opposing positions, personalisation becomes necessary. The guide responds with per-category filters, collapsible progressive disclosure, opt-in toggles and multi-format output.

External constraints reshaped delivery in two ways.

Multi-format publishing from a single draft was unavoidable in the PTSA pilot. Procurement could not accept content served from a third-party site inside the pilot window, and asked for a Word document and a PDF for distribution through the accessibility resources suite.

Several originally scoped dynamic data feeds were unavailable on the Adelaide Metro network or at any comparable Australian transit hub. Operator-authored content through the admin tool was built in their place.

11.2 Outstanding questions

- **Severity wording.** No Australian standard exists for sensory-severity terminology, and Aspect's preferred wording was pending at the time of writing. The platform supports per-venue overrides as an interim measure.
- **Transfer beyond this cohort and venue.** The framework's six components were derived from research with Autistic lived-experience participants at Adelaide Railway Station. Whether the same components surface for other sensory-sensitive constituencies and at other venue types is the next-stage research question.
- **Standards gap beyond severity wording.** No Australian standards exist for sensory-mapping colour scales, severity terminology or icon sets. Three certification frameworks (Aspect's Autism Friendly Accreditation, the Autism SA Autism Friendly Charter and Dementia Australia's design principles) operate without cross-recognition.
- **Scale of the evidence base.** Milestone 3 is small-scale qualitative: six lived-experience participants across the two focus groups, plus three Aspect auditors on the admin tool.

  Recruitment via Aspect concentrated the lived-experience cohort on Autistic participants. The pilot is a single venue. The recommendations that would address these limitations are set out in §10.

11.3 Acknowledgements

This research is funded by iMOVE CRC and supported by the Cooperative Research Centres Programme, an Australian Government initiative.

The research team thanks Autism Spectrum Australia (Aspect) and its Autism Friendly team for partnership across the milestone: recruitment, co-designed question sets, facilitation of both focus groups, and review between iterations. Thanks also to PTSA for hosting the pilot at Adelaide Railway Station, and to the focus-group participants whose lived-experience insight shaped the design.

Appendices

Appendix A. Milestone 2 reuse note

The Milestone 2 deliverable (Static Sensory Mapping Report, 2025) carries three items reused in Milestone 3: the Aspect Sensory Sensitivity Audit for Adelaide Railway Station and its annotated site maps; the mapping methodology and sensory trigger categories; and the validation focus-group feedback and its findings. The audit itself sits with the Milestone 2 deliverable and is not reproduced here.

Findings carried forward are summarised in §2.3. The map-to-guide pivot is documented in §3 and §5.

Appendix B. Audit and quality summary

Three automated reviews of the Milestone 3 build were undertaken on 19 May 2026: a security review, an accessibility review against WCAG 2.2 AA, and a build-quality and continuous-integration review. Full audits are on file with the project team.

B.1 Security review

Scope: authentication and authorisation, input validation, database and file-storage access rules, client-side injection paths, secrets and configuration, and dependency security. The auditor approved continued deployment on condition the two high-severity findings are remediated.

Severity	Count
Critical	0
High	2
Medium	4
Low	4

> *‘Conditionally approved for continued production deployment subject to remediation of the two HIGH findings.’ (Automated security review, 19 May 2026)*

The two high-severity findings sit within the authorisation layer and are fixable in under 30 lines of code per the auditor. The broader-deployment audits required are an independent third-party penetration test and a formal threat model of the AI content pipeline (§10.4).

B.2 Accessibility review

Scope: the public guide, the admin tool, the audio output, and the design-system colour tokens, audited against WCAG 2.2 AA. The auditor confirmed that none of the findings block release.

Severity	Count
Critical	0
High	0
Medium	5
Low	8
Informational	2

> *‘None of the findings block release.’ (Automated WCAG 2.2 AA review, 19 May 2026)*

The five medium findings are all colour-contrast or motion items, each with a planned mitigation. The broader-deployment audits required are an independent third-party WCAG sign-off with real-user assistive-technology session testing, and an export-artefact audit covering the PDF, Word and audio outputs (§10.4).

B.3 Build quality and continuous integration

Every push to the production branch runs five automated quality jobs. The deploy step is blocked unless every job passes.

Job	What it checks	Blocks deploy
Lint, test, build	Code style, unit tests, application build	Yes
Dependency security (web)	Third-party package advisories at high severity or above	Yes
Dependency security (server)	Third-party package advisories at high severity or above	Yes
Accessibility	Automated WCAG 2.2 AA	Yes

Job	What it checks	Blocks deploy
	check against the public homepage	
Performance and accessibility score	Synthetic-lab score thresholds (accessibility ≥ 0.95, performance ≥ 0.80)	Yes

Public-viewer initial load on the 27 May 2026 build measured approximately 144 KB compressed:

Chunk	Raw KB	Gzip KB
Application code (home page, guide page, public layout, shared, data-fetching)	369	116
User-interface framework (React, Router)	48	17
Styling	59	11
Public-viewer initial total	**476**	**144**

The authentication and editing-database SDK (492 KB raw, 150 KB compressed) is no longer part of the public-viewer initial load. It loads on demand only when an operator signs in to the admin tool, or when a public visitor requests a guide's spoken-audio version. Other large admin-only resources, including PDF and document generation, behave the same way. Total uncompressed application size across all chunks is 3.5 MB. A representative-mobile time-to-interactive measurement on the production build is one of the numbers to capture before broader deployment (§10.4).

Appendix C. Glossary

Term	Definition
Aspect	Autism Spectrum Australia, the national peak

Term	Definition
	body. Aspect's Autism Friendly team recruited participants, co-designed the question sets, ran the sessions and reviewed each iteration cycle. Aspect facilitators hold both professional roles and lived-experience perspective.
Sensory Guide	The Milestone 3 web product: a published per-venue planning-oriented guide describing the sensory experience of each area, with category filters, severity badges, photographs, an on-demand map and multi-format outputs (web, embeddable partner version, PDF, Word, audio).
Action plan	Milestone 2's term for action-oriented content describing what to expect and how to prepare for a venue. Carried into Milestone 3 as the form the Sensory Guide takes.
BindiMaps	The accessible-wayfinding platform on which the Sensory Guide is built. Provides venue map data, authentication and distribution channels (app, kiosk, web embed).
Sensory category	A sensory dimension the guide describes per area, and the filter axis on the public guide. Nine are deployed at Adelaide Railway Station (sound, light, smell, crowds, temperature, touch, movement, vibration, air quality); the framework names five (sound, light, smell, crowds, temperature) as the focus-group-tested core (§9.1.3).
Severity	The intensity of a sensory category in an area. The built product uses Low, Medium or High with per-venue overridable labels (§4.4.7); the framework names Low, Moderate or High (§9.1.3). Both forms appear in this report.

Term	Definition
Quiet Room	A designated low-stimulus space inside a venue. Aspect runs an accreditation scheme for certified Quiet Rooms specifically, distinct from venue-wide accreditation.
GTFS-RT	General Transit Feed Specification, Real-Time variant. The open standard for live transit data. The `OccupancyStatus` field, structurally available in the protocol, is the field Adelaide Metro does not populate (§8).
WCAG 2.2 AA	Web Content Accessibility Guidelines 2.2, level AA. The accessibility standard the Milestone 3 build was audited against (§7.5, Appendix B).
Place-subject-detail	The Sensory-Sensitive Wayfinding Framework's content hierarchy: venue (place), area (subject), per-category sensory detail (detail). See §9.1.1.
Before / during / after	The Sensory-Sensitive Wayfinding Framework's user-phase frame: planning (before), in-venue use (during), feedback loop (after). See §9.1.2.
AI-as-suggester	The pattern in which AI proposes specific sentence-level edits with rationale and the human accepts or dismisses each one. Inverts the original AI-as-author pipeline (§5.4.1, §9.1.5).
Multi-format output	Five published artefacts from one source draft: interactive web, embeddable partner version, PDF with QR codes, Word document, audio MP3 (§4.4.3, §9.1.4).
Public Transport South Australia (PTSA)	The customer-facing public-transport authority under the South Australian Government's Department for Infrastructure

Term	Definition
	and Transport. Pilot partner. Its Customer Strategy and Inclusion Lead was the partner contact.
Adelaide Metro	The wider South Australian public-transit network of which Adelaide Railway Station is part.
FG1 / FG2	Focus Group 1 (24 March 2026) and Focus Group 2 (19 May 2026). The two Milestone 3 user-research sessions (§3).

Appendix D. References and prior-art list

D.1 Bibliography

Inclusive-design canon underpinning §2:

- Hamraie, A. (2017). *Building Access: Universal Design and the Politics of Disability.* University of Minnesota Press. https://www.upress.umn.edu/9781452955568/building-access/. Canonical treatment of the curb-cut effect.
- Mace, R., Center for Universal Design (1997). *The Principles of Universal Design, Version 2.0.* North Carolina State University. https://design.ncsu.edu/wp-content/uploads/2022/11/principles-of-universal-design.pdf.
- Microsoft (2016). *Inclusive Design Toolkit Manual.* https://download.microsoft.com/download/b/0/d/b0d4bf87-09ce-4417-8f28-d60703d672ed/inclusive_toolkit_manual_final.pdf. Persona spectrum and situational-disability framing.
- Oliver, M. (1990). *The Politics of Disablement.* Macmillan. Articulation of the social model of disability.
- Union of the Physically Impaired Against Segregation (UPIAS) (1976). *Fundamental Principles of Disability.* https://disability-studies.leeds.ac.uk/wp-content/uploads/sites/40/library/UPIAS-fundamental-principles.pdf. The originating social-model statement cited at §2.1.

Foundational sensory-sensitivity background cited in §1:

- American Psychiatric Association (2013). *Diagnostic and Statistical Manual of Mental Disorders*, 5th ed. (DSM-5). Washington, DC: American Psychiatric Publishing. https://doi.org/10.1176/appi.books.9780890425596. Autism diagnostic criterion B.4 lists

'hyper- or hyporeactivity to sensory input or unusual interest in sensory aspects of the environment'.

- Dean, E. E., Little, L., Tomchek, S., Wallisch, A., and Dunn, W. (2022). 'Prevalence models to support participation: sensory patterns as a feature of all children's humanity.' *Frontiers in Psychology*, 13: 875972. https://doi.org/10.3389/fpsyg.2022.875972.
- Nolan, C., Doyle, J. K., Lewis, K., and Treanor, D. (2023). 'Disabled Students' perception of the sensory aspects of the learning and social environments within one Higher Education Institution.' *British Journal of Occupational Therapy*, 86(5): 367–375. https://doi.org/10.1177/03080226221126895.
- Tomchek, S. D., and Dunn, W. (2007). 'Sensory Processing in Children with and without Autism: A Comparative Study Using the Short Sensory Profile.' *American Journal of Occupational Therapy*, 61(2): 190–200. https://pubmed.ncbi.nlm.nih.gov/17436841/.

Peer-reviewed empirical work cited in §2:

- Haas, K., Wilson, N. J., Cordier, R., Vaz, S., and Chung-yeung Lee, H. (2020). *The experiences of young autistic adults in using metropolitan public transport.* Brisbane: Cooperative Research Centre for Living with Autism (Autism CRC). https://www.autismcrc.com.au/sites/default/files/reports/3-041_Young-autistic-adults-on-public-transport_Report_2020.pdf.
- MacLennan, K., O'Brien, S., and Tavassoli, T. (2021). 'In Our Own Words: The Complex Sensory Experiences of Autistic Adults.' *Journal of Autism and Developmental Disorders.* https://pmc.ncbi.nlm.nih.gov/articles/PMC9213348/.
- Neil, L., Olsson, N. C., and Pellicano, E. (2016). 'The Relationship Between Intolerance of Uncertainty, Sensory Sensitivities, and Anxiety in Autistic and Typically Developing Children.' *Journal of Autism and Developmental Disorders.* https://pmc.ncbi.nlm.nih.gov/articles/PMC4860201/.

Project deliverables:

- *Static Sensory Mapping Report (Milestone 2).* 2025. Source of the structural conventions reused in this report (Appendix A); §9.3 quote on action-plan preference cited per §5.2.
- Aspect Sensory Sensitivity Audit, Adelaide Railway Station. Carried forward from Milestone 2; §2.1 quote on the escalator clunk cited per the audit.

Aspect organisational material (accessed 19 May 2026):

- Aspect, *Access in the arts: What makes a space autism-friendly* (blog). https://www.aspect.org.au/blog/access-in-the-arts-what-makes-a-space-autism-friendly.
- Aspect, *Aspect Autism Friendly Accreditation.* https://www.aspect.org.au/our-services/autism-friendly/aspect-autism-friendly-accreditation. Quiet Room accreditation criteria including the feedback commitment cited at §9.1.2.

Independent corroborating evidence:

- Buro Happold (2026). *Sensory Mapping: The design and use of sensory maps for the built environment.* Whitepaper. https://www.burohappold.com/wp-content/uploads/2026/03/BD302_Sensory-Mapping-The-design-and-use-of-sensory-maps-for-the-built-environment_01-1.pdf. 280-respondent survey cited at §3, p. 5 and p. 16.

Policy:

- Australian Government (2025). *National Autism Strategy 2025–2031.* https://www.health.gov.au/sites/default/files/2025-07/national-autism-strategy-2025-2031.pdf. Cited at §10.2 for the sensory-friendly health undertaking.

D.2 Comparable prior art: sensory-sensitive wayfinding guides, social stories, action plans

Distilled from `02a-prior-art.md` Section A. Each entry: title, organisation, format, URL.

Museums and galleries

- *Visual Story.* Australian Museum, Australia. Web. https://australian.museum/visit/access/visual-story/
- *Sensory Friendly Resources* (4 PDFs, age-segmented). WA Maritime Museum, Australia. https://visit.museum.wa.gov.au/maritime/sensory-friendly-resources
- *Information for neurodiverse visitors.* NGV, Australia. Web hub plus per-exhibition PDFs. https://www.ngv.vic.gov.au/plan-your-visit/access/information-for-neurodiverse-visitors-or-those-with-sensory-sensitivities/
- *Visual Story.* Tate Modern, UK. Web. https://www.tate.org.uk/visit/tate-modern/visual-story
- *Disability access* (multi-format programme; audio plus tactile-with-QR). V&A, UK. https://www.vam.ac.uk/info/disability-access
- *Morning at the Museum* (multi-institution programme). Smithsonian, US.

Transport

- *Social Stories.* Hobart Airport, Australia. PDF. https://hobartairport.com.au/travellers/airport-guide/accessibility/social-stories/
- *Sydney Trains and Sydney Metro visual stories* (Hidden Disabilities Sunflower, Aspect collab). Transport for NSW, Australia. https://transportnsw.info/news/2024/extra-support-for-hidden-disabilities
- *Hidden Disabilities* (PDF stories plus YouTube walkthroughs). Cairns Airport, Australia. https://www.cairnsairport.com.au/travelling/airport-guide/hiddis/

- *Hidden Disabilities* (sensory map plus AAC communication boards). Townsville Airport, Australia. https://www.townsvilleairport.com.au/accessibility/accessibility-overview/hidden-disabilities
- *Hidden Disability Program* (three-tier sensory zoning across 4 terminals). Melbourne Airport, Australia. https://www.melbourneairport.com.au/hidden-disability-program
- *Hidden Disabilities* (4 directional sensory stories). Canberra Airport, Australia. 7 PDFs. https://www.canberraairport.com.au/passenger-guide/help-services/hidden-disabilities

Shopping centres

- *Sensory Shopping Guide.* Stockland Burleigh Heads, Australia. Web plus PDF map. https://www.stockland.com.au/shopping-centres/centres/stockland-burleigh-heads/offers-and-events/news/sensory-shopping-guide

Arts and performing arts

- *Visual Story Library for Relaxed Performances.* Arts Centre Melbourne, Australia. Web library. https://www.artscentremelbourne.com.au/visit/accessibility/when-you-visit/relaxed-performances/visual-story-library
- *Sensory Friendly Space* (per-event-type visual stories). MCG, Australia. PDFs. https://www.mcg.org.au/plan-a-visit/accessibility-inclusion/sensory-friendly-space

National parks and outdoor recreation

- *Autism-Friendly Visits* (10 parks; PDF, editable DOCX, Key Word Sign Australia versions). Parks Victoria, Australia. https://www.parks.vic.gov.au/get-into-nature/all-abilities-access/autism-friendly-visits
- *Accessibility* (sensory maps across 4 properties, social scripts across 3; Amaze collab). Zoos Victoria, Australia. https://www.zoo.org.au/accessibility

Healthcare

No verified entries. Hospital passport references exist at Sydney Children's Hospitals Network and NSW Health, but no actual passport or pre-visit resource met the verification bar, a healthcare gap taken up in §10.2.

D.3 Organisations engaging with sensory accessibility (market-breadth evidence)

Distilled from `02a-prior-art.md` Section B. Mostly sensory **maps** (the approach the research rejected as the primary shape), retained as evidence of cross-sector interest in sensory accessibility. Supports §10's broader-deployment argument.

Museums and galleries (sensory maps)

- Royal Academy (UK), Science Museum London (UK), Wallace Collection (UK), MCA Australia, National Museum of Australia (Mr Squiggle), Australian Museum (visitor map), Australian Sports Museum, Scienceworks, Melbourne Museum (2022 and 2019 versions). URLs in source evidence pack.

Transport (sensory maps)

- Melbourne Airport (maps hub), Adelaide Airport (Terminal Sensory Map), Gold Coast Airport (Terminal Sensory Map).

Shopping centres (sensory maps)

- Stockland Burleigh Heads (A3 sensory map), Stockland Point Cook, Stockland Rockhampton, DFO Brisbane, Burwood Brickworks.

Local government, libraries and civic (sensory maps)

- Monash Public Library Service (web hub plus multi-branch PDF), NC State University Libraries (US, including a lighting-specific map), Georges River Council.

Arts, sports and events (sensory maps)

- Australian Open (Sensory Map 2024).

Education and research (sensory maps)

- University of Wollongong (Sensory Discovery Space Map), Project ENABLE (Syracuse University, US, professional development for librarians and teachers).

Guidance and how-to references

- Historic England, Sensory Trust (UK), Send in Museums (UK), Australian Retailers Association (neurodiverse-shopper guidance), Aspect Sensory Maps (Wetherill Park).

Standards, certification and peak-body guidance

- Aspect Autism Friendly Accreditation (Autism Spectrum Australia). Aspect's accreditation applies specifically to certified Quiet Rooms, not to venues at large. https://www.aspect.org.au/our-services/autism-friendly/aspect-autism-friendly-accreditation
- Autism Friendly Charter (Autism SA). Training-led certification, public business directory, commitment badges. https://autismfriendlycharter.org.au/
- Dementia Australia, *Designing dementia-friendly care environments.* Design principles for dementia-friendly built environments. https://www.dementia.org.au/professionals/designing-dementia-friendly-care-environments

Appendix E. Demonstration videos

Three short screen recordings show the Milestone 3 build in use: the public guide on desktop and on mobile, and the operator workflow from import to publish. Each is hosted on Vimeo; follow the link or scan the QR code. Video cannot be embedded in a Word document, so the recordings are linked rather than played inline.

E.1 Planning a trip on the public guide (desktop, 2:37)

A visitor prepares for a visit on the public Sensory Guide: moving through the Before, During and After structure, filtering by sensory category, and expanding and collapsing area sections. https://vimeo.com/1197919097

E.2 Using the guide in the venue (mobile, 1:13)

The same guide on a phone, as a visitor would use it on arrival at the venue. https://vimeo.com/1197919096

E.3 Operator workflow from import to publish (4:44)

A venue operator builds a guide end to end: importing a source PDF, editing content (including dictating by voice), and publishing the live guide. https://vimeo.com/1197919095