LIFT Study 01

OHCA Drone-Delivered AED Feasibility Research Proposal

A feasibility assessment of aerial AED delivery for out-of-hospital cardiac arrest in New South Wales.

***No Kill Switch Research Programme** — 2026*

Public consultation draft v0.1.0.

PUBLICATION

LIFT Study 01: OHCA Drone-Delivered AED Feasibility (NSW). Version 0.1.0, 2026. No Kill Switch Research Programme.

SUGGESTED CITATION

No Kill Switch Research Programme. LIFT Study 01: OHCA Drone-Delivered AED Feasibility. Sydney: No Kill Switch, 2026. Available at lift.nokillswitch.com.

COPYRIGHT

CHAPTER 01

Introduction

Preface

I’ll never forget the first time I was alerted as a GoodSAM responder. During my final year of Paramedicine study, I was due to go on shift with NSW Ambulance at 11am during my penultimate placement.

The alert came through at 07:15 on a Thursday morning and I was on scene three minutes later at a house on the corner of a main road, opposite a service station. The patient, an elderly grandmother, had collapsed in a narrow passage between a downstairs bedroom and the toilet where she had collapsed. Her granddaughter was doing CPR, but the narrow passage made it ineffective. The family were distressed and CPR was being directed by the 000 call taker on speaker. I took over the scene, confirmed arrest, and with help moved the patient to get proper access. I administered an airway, and asked the teenage grandson to check for an AED at the service station despite one not on the GoodSam AED map. He came back empty-handed.

There were five AEDs within 1.5 kilometres of us: one at a church, at a GP surgery, another church, a primary school, and a community centre but none were accessible. None of the organisations were due to open for at least 40 minutes. Even if they were open, or the AEDs accessible, the round trip time especially in morning traffic, was too long to make a difference.

So we worked the arrest with what we had; CPR, an airway adjunct, a bag-valve mask, and untrained bystanders, until the ambulance arrived 15 minutes after me. We never achieved ROSC.

What struck me was the absence of the one intervention most likely to change the outcome: defibrillation, not any failure of clinical care.

A few months later, I was alerted by GoodSam again. This time at 1am. This patient lived about 600 metres from my house and I arrived just before the ambulance. Again, there were AEDs nearby and again, none were accessible at that time of night.

This patient, however, had a shockable rhythm, and despite rapid response across empty streets, we didn’t get a defibrillator onto him until 15 minutes after his arrest. Even so, we achieved ROSC, but ultimately he died in the ED.

These experiences changed how I think about cardiac arrest. Despite my clinical training, despite my passion in cardiology and research in all the amazing clinical interventions we wield, BLS, ALS, retrieval and eCPR, it strikes me that this is a system design problem.

We have trained responders, established protocols and proven technology. We simply don’t have alignment between where cardiac arrests occur and where defibrillators are available. Surely the biggest impact we can make to survivability of OHCA is tweaks to what happens after the ambulance arrives, but how quickly we can defibrillate and get high quality CPR happening?

> *For every minute that passes without defibrillation, the chance of survival from cardiac arrest decreases by 7 to 10 percent.*
>
> — American Heart Association, *Early Defibrillation Advisory Statement* [1]

A note on data labels

Every quantitative claim in this report carries one of four provenance labels so the reader can see exactly what kind of signal is being used. The table below names the four labels and the treatment each implies.

Label	Meaning
OBSERVED	Figure taken directly from a published or curated source.
ESTIMATED	Planning-grade scenario output; bounded, explicit assumptions.
SYNTHETIC	Demographic proxy; pending verification against NSW Ambulance data.
INFERRED	Candidate geometry derived from SA2 centroids; not dispatch-grade.

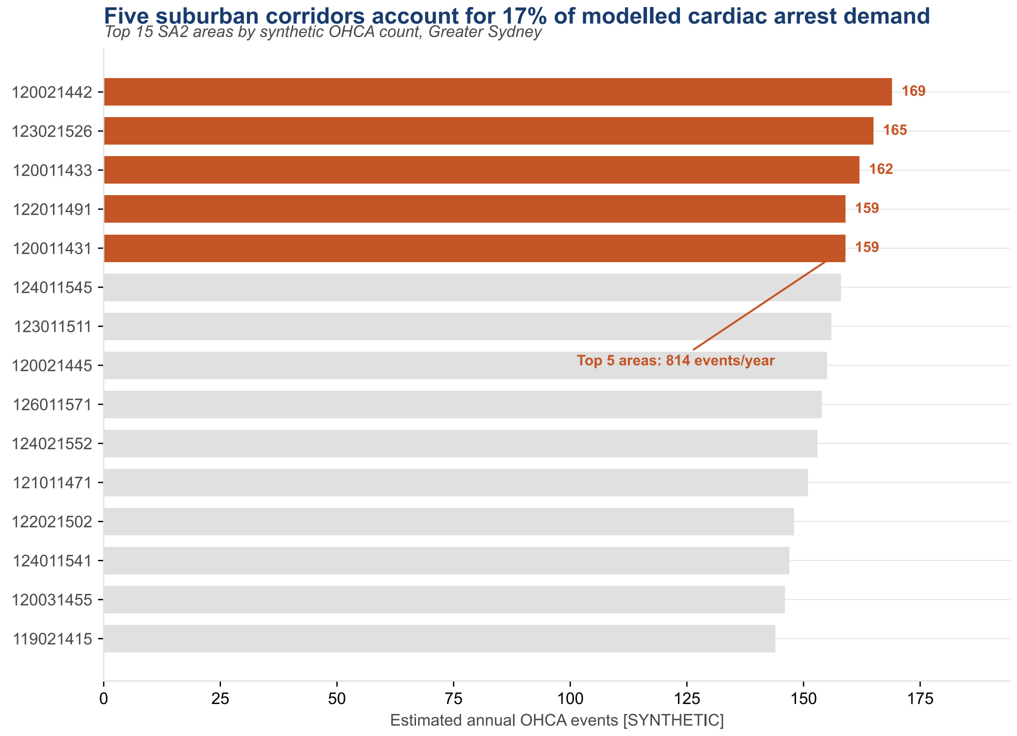
Five suburban corridors account for 17% of modelled cardiac arrest demand

Top 15 SA2 areas by synthetic OHCA count, Greater Sydney

Top 5 areas: 814 events/year

Data label: synthetic/estimated. Source: synthetic_ohca_hotspots.csv

1 Figure 1.1 — The SEIFA-weighted synthetic OHCA hotspot distribution across 37 Greater Sydney SA2s. Each bar is one SA2, ordered by synthetic annual count. The shape of the curve, a long tail dominated by a south-western cluster, motivates the demand-signal chapter that follows. Data label: synthetic.

CHAPTER 02

Executive Summary

Thesis. A drone-delivered AED network could meaningfully improve out-of-hospital cardiac arrest survival across NSW by closing the residential response-time gap that fixed public-access AEDs cannot reach. The rural and regional case is arithmetically stronger than metro alone, because the ambulance response-time deficit is larger at those tiers and the survival uplift from a 4-minute drone is therefore more leveraged. A funded 12-month three-site, three-tier pilot (Metro / Regional / Rural NSW) is the appropriate next step to test whether the modelled benefits hold under Australian operational conditions.

Confidence level. Feasibility-stage. The time-survival relationship is well established; the operational, regulatory, and bystander-retrieval variables are not yet measured in Australian conditions at any of the three tiers.

Three key findings.

1. Across the three-tier pilot the model estimates **21 to 86 additional lives per year** within the covered population, depending on bystander-retrieval rate (100% retrieval down to 30%) and decay-parameter sensitivity (±20% on the Valenzuela-Larsen survival curve). The Rural per-event case is the strongest (Cat-1 response deficit more than doubles the Metro deficit); the Metro tier delivers the largest absolute volume.
2. The hardware-loaded 5-year envelope for a five-base network is **$1.5M; the fully-loaded operating envelope across the three-site pilot is $3–5M base** ($2M optimistic through $8.5M pessimistic) once CASA approval, dispatch integration, staffing, insurance, evaluation, community engagement, and tier-varied regulatory work are priced (Chapter 09 BOM).
3. Cost per QALY spans **~$450 to $5,500/QALY** across the operating-model envelope, an order of magnitude below published public-access AED benchmarks of $30,000–$55,000/QALY at every point in the range [2].

Three unresolved risks.

1. CASA approval for medical-delivery BVLOS over populated areas is case-by-case and has no Australian precedent at this specific payload class. Tier-stratified: the Metro pilot site is the hardest submission; Rural is simpler because the airspace is uncontested and the SORA-AU ground-risk class drops (Chapter 11).
2. Bystander AED retrieval at scale has been demonstrated once globally (Karolinska/Everdrone 2021). Australian rates are unknown and may differ between Metro, Regional, and Rural bystander populations — the three-tier pilot measures this directly.
3. NSW Ambulance dispatch integration across three sites is a non-trivial negotiation rather than a technical fact; station-based siting at each tier depends on a formal co-location memorandum of understanding.

The specific pilot ask. A funded 12-month three-site pilot at **Campbelltown (Sydney metro) + Port Macquarie (Regional) + Far West / Broken Hill (Rural)**, estimated at **~$3–5M all-in**, delivering tier-stratified response-time comparison against ambulance, bystander retrieval rates, and preliminary clinical-outcome data to inform the scaled deployment decision (Chapter 12).

Download the briefing for your audience.

- Services Proposal (PDF) — for NSW Ambulance, NSW Health, NSW Rural Health, and government readers
- Investment Feasibility (PDF) — for capital partners and unit-economics readers
- Research Proposal (PDF) — for academic collaborators and grant reviewers

These caveats constrain the claim but do not overturn the survival-time evidence base. The relationship between time-to-defibrillation and survival from shockable cardiac arrest is among the most replicated findings in emergency medicine: faster defibrillation saves lives. The open question is whether drones can reliably deliver that speed advantage in Australian conditions across Metro, Regional, and Rural NSW, at a fully-loaded cost the sponsor can absorb, under a regulatory posture CASA will accept. A funded three-tier pilot is the only credible way to find out.

CHAPTER 03

Spatial distribution of cardiac arrest in NSW

A defibrillator-carrying drone network is, at heart, a spatial optimisation problem: where the events occur dictates where the bases must be, and how far a drone must fly to reach each new collapse. This chapter characterises the spatial distribution of out-of-hospital cardiac arrest (OHCA) across NSW. It uses three observed sources — BHI Healthcare Quarterly ambulance activity and response times at SA3, the NSW Ambulance Cardiac Arrest Registry state-aggregate outcomes, and the ABS ASGS 2021 Remoteness Area classification — alongside the retained modelled demand allocation score documented in Appendix M.

Three-tier framing. Candidate pilot sites are stratified into three tiers using the ABS Remoteness Area classification. **Metro** covers Major Cities of Australia, with Sydney, Newcastle, and Wollongong ranked as separate metros. **Regional** covers Inner Regional and Outer Regional Australia. **Rural** covers Remote and Very Remote Australia. The statewide three-tier candidate ranking names a top-1 headline plus top-2 secondary candidate per tier, with the metros split three ways; the ranking dataset is carried in the reproducibility appendix, and the final set will be re-run when restricted NSW Ambulance registry data is received.

The Greater Sydney SA2 analysis described below is retained as the Metro(Sydney) view, and is now one of three Metro read-outs rather than the whole picture. Two complementary views are shown below: a rate choropleth that normalises for population (task 3.3), and a proportional-symbol map that reports absolute event volume (task 3.4). Read together, they reveal a pattern that is not obvious from either map alone.

Synthetic-cohort caveat. The figures and rankings below descend from a proxy dataset that combines the ABS 2021 Greater Sydney SA2 population and age structure with the ABS 2021 SEIFA Index of Relative Socio-economic Disadvantage (IRSD), anchored to a Hasselqvist-Ax 2015 incidence base rate. The derivation is documented in full in Appendix M: Methodology: synthetic OHCA proxy with SEIFA IRSD weighting. Treat every absolute count on this page as pending verification against NSW Ambulance incident data; relative ordering between SA2s is expected to hold under replacement with the real cohort.

Where rates concentrate

OHCA rates concentrate in Sydney's disadvantaged south-west once SEIFA IRSD is weighted in

Synthetic OHCA incidence per 100,000 population · 5-class quantile · 37 SA2s, Greater Sydney · proxy pending NSW Ambulance data

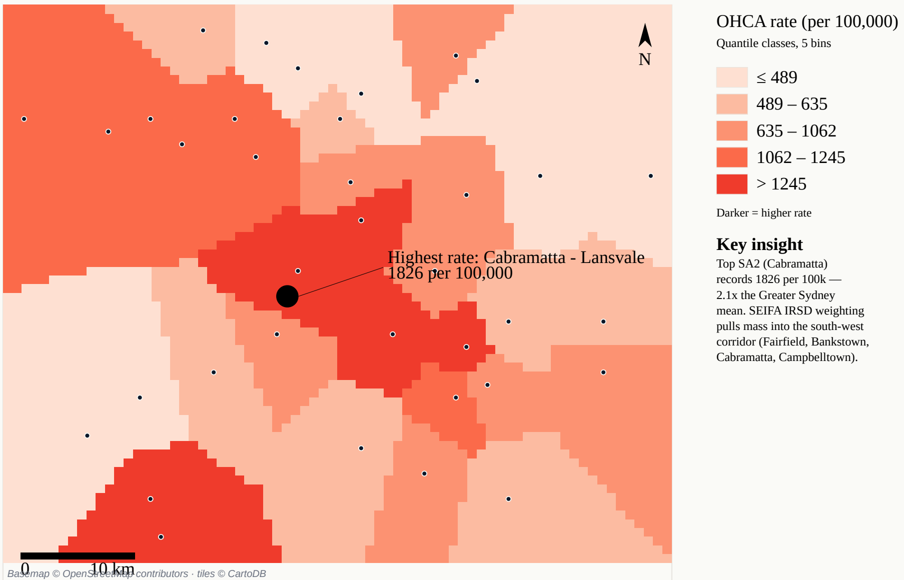

Key insight

Top SA2 (Cabramatta) records 1826 per 100k — 2.1x the Greater Sydney mean. SEIFA IRSD weighting pulls mass into the south-west corridor (Fairfield, Bankstown, Cabramatta, Campbelltown).

Source: synthetic OHCA hotspots (data/derived/synthetic_ohca_hotspots.csv), derived from ABS SA2 2021 geography (data/curated/sa2_greate
Areal unit: ABS Statistical Area Level 2 (SA2), Greater Sydney (n=37). MAUP caveat: spatial pattern depends on the choice of SA2 boundaries
Projection: equirectangular, scaled at the latitude midpoint. Choropleth surface: nearest-centroid (Voronoi) raster proxy rendered at 72x64 cells -
Small-numbers caveat: all SA2s carry ≥ 33 synthetic cases, above the suppression threshold; Empirical Bayes smoothing remains advisable at fi
Figure by Grant (mapping agent), LIFT Study 01 · Sprint 20 · re-rendered Sprint 26 (SEIFA IRSD weighting) · 2026-04-21.

Figure 3.1 — SA2 demand allocation score (SEIFA-IRSD-weighted), ABS SA2 geography, 5-class quantile, rendered over a CartoDB Positron Sydney base map (© OpenStreetMap contributors, CC BY). Surface rendered as a nearest-centroid (Voronoi) raster proxy because ABS SA2 polygon boundaries are not shipped with this repository; see the map footnote for full provenance, the MAUP caveat, and the allocation-vs-incidence distinction documented in Appendix M.

Key insight. Rate concentration is a joint geography of age *and* disadvantage. Under the SEIFA-IRSD-weighted proxy (Appendix M), the highest synthetic rates sit in Sydney’s south-western corridor: Fairfield, Cabramatta–Lansvale, Bankstown, Campbelltown-South and Campbelltown-North all land in the top quantile. Age structure still matters, but the IRSD weighting pulls the rate surface decisively away from the leafy northern-shore band the earlier age-only proxy highlighted. The policy implication is inverted from that earlier reading: a “wealthy boomers” heuristic would mis-site base placements away from the SA2s the combined age-plus-deprivation signal actually prioritises. The choropleth also carries the full MAUP caveat: re-aggregating the same cohort to SA3 or LGA boundaries would smooth the south-west cluster into a single mid-range polygon, and pooling across LGA would re-introduce the wealthy-suburb age-signal artefact.

Where the volume sits

Fairfield / Bankstown / Cabramatta top the SEIFA-weighted synthetic OHCA proxy

Proportional circle (area ∝ annual cases) · 37 SA2 centroids · Greater Sydney synthetic cohort · proxy pending NSW Ambulance data

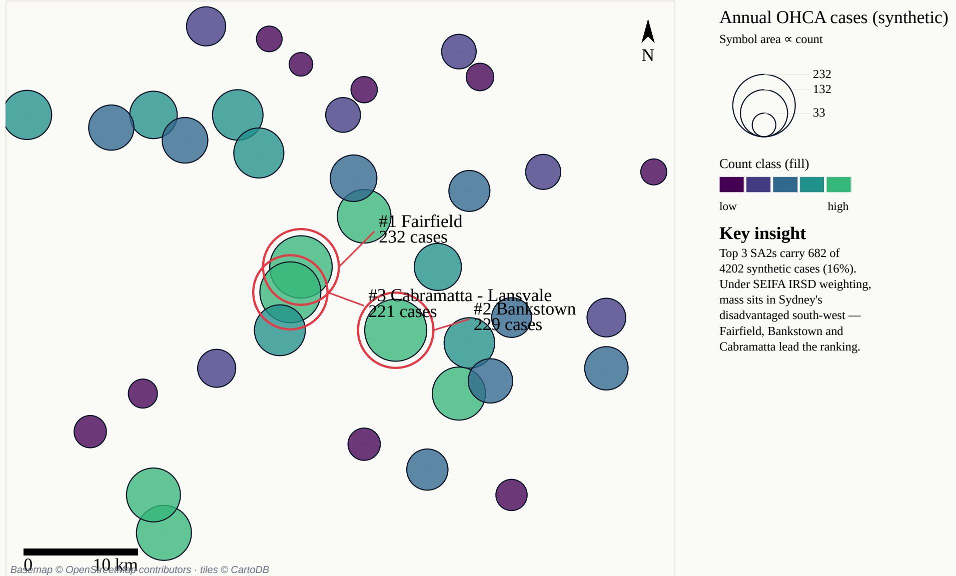

Key insight

Top 3 SA2s carry 682 of 4202 synthetic cases (16%). Under SEIFA IRSD weighting, mass sits in Sydney's disadvantaged south-west — Fairfield, Bankstown and Cabramatta lead the ranking.

Source: data/derived/synthetic_ohca_hotspots.csv (synthetic annual counts) · SA2 centroids from data/curated/sa2_greater_sydney.csv (ABS 20
Areal unit: ABS Statistical Area Level 2 (SA2) centroid (n=37), Greater Sydney. MAUP note: counts aggregate all within-SA2 events to a single
Projection: equirectangular, scaled at the latitude midpoint. Symbol area is linearly proportional to count (no Flannery perceptual correction app
Small-numbers caveat: smallest symbol represents 33 synthetic cases; all symbols exceed the n<5 suppression threshold.
Figure by Grant (mapping agent), LIFT Study 01 · Sprint 20 · re-rendered Sprint 26 (SEIFA IRSD weighting) · 2026-04-21.

Figure 3.2 — Proportional-symbol map of synthetic annual OHCA counts by SA2 centroid, Greater Sydney (n = 37), rendered over a CartoDB Positron Sydney base map (© OpenStreetMap contributors, CC BY). Symbol area is linearly proportional to count; fill colour (Viridis) repeats the magnitude encoding for redundancy and colourblind safety.

Key insight. Volume and rate now point at the same part of the city, the disadvantaged south-west. The top-three SA2s by synthetic volume (Fairfield, Bankstown, and Cabramatta–Lansvale) between them carry approximately 16% of synthetic cases, and the top-five extends this cluster south to Campbelltown-South and Campbelltown-North. Every SA2 in the top-five sits in SEIFA IRSD decile 1 or 2 (most disadvantaged), confirming that the combined age-plus-IRSD proxy is the logical first-pass driver of drone-base placement. Note the deliberate encoding: a choropleth would misrepresent this view because event *counts* (as opposed to rates) are an extensive variable and must not be mapped by area fill. The proportional symbol map is the cartographically honest form for the question “where are the biggest hotspots, in absolute terms?”

Candidate bases: ambulance stations versus synthetic SA2 centroids

The set-cover coverage model in chapter 07 relies on a candidate pool of possible base locations. The original pool was synthetic SA2 centroids anchored on the top-10 SEIFA-weighted hotspots. That pool is defensible as a first-pass modelling exercise but it invites the obvious counter-question from a reader at NSW Ambulance: why not base drones at existing ambulance stations? Co-location reuses electrical supply, security, and 24/7 crewed presence, and it simplifies the dispatch-integration story since the drone is launching from the same site the paramedic crew already works from.

A hand-curated pool of 44 NSW Ambulance stations inside the Greater Sydney bounding box provides a first-class candidate pool, hand-curated from the NSW Ambulance public station directory cross-referenced against OpenStreetMap `amenity=ambulance_station` records. The file sits at the NSW Ambulance station reference dataset (reproducibility appendix); licence and source attribution are carried in the CSV header.

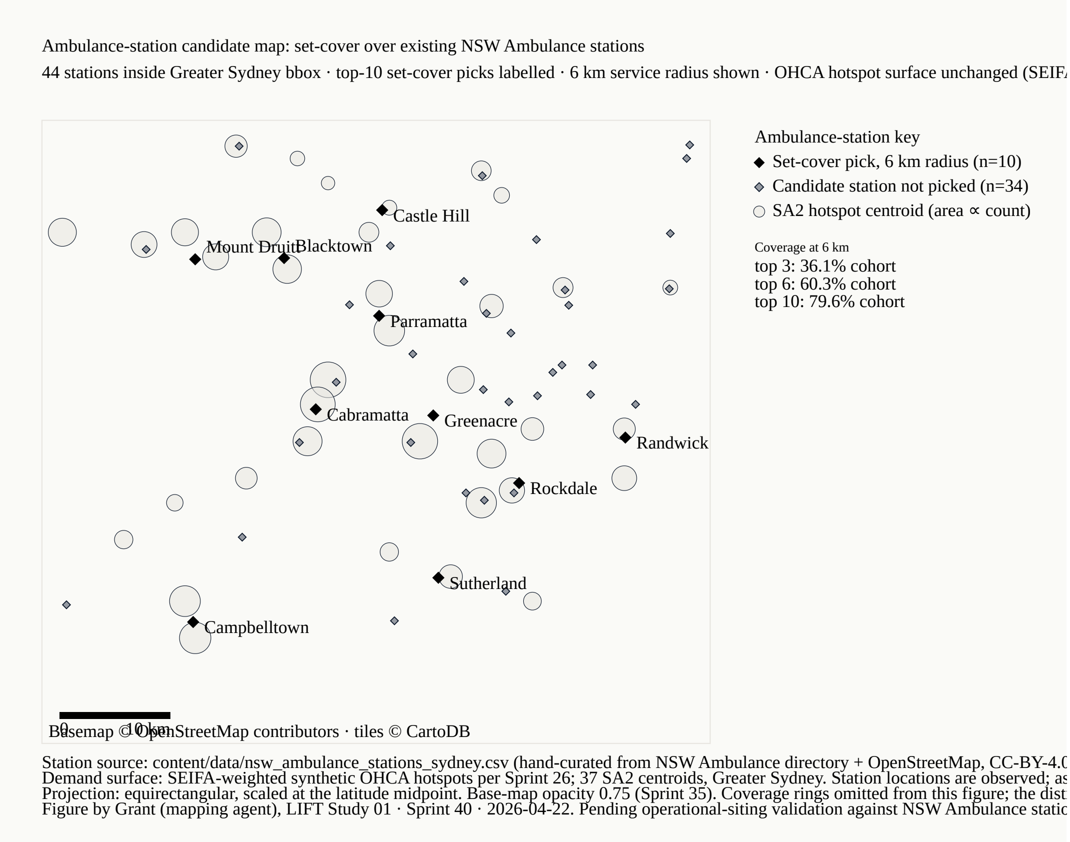

Ambulance-station candidate map: set-cover over existing NSW Ambulance stations

44 stations inside Greater Sydney bbox · top-10 set-cover picks labelled · 6 km service radius shown · OHCA hotspot surface unchanged (SEIFA

Station source: content/data/nsw_ambulance_stations_sydney.csv (hand-curated from NSW Ambulance directory + OpenStreetMap, CC-BY-4.0
Demand surface: SEIFA-weighted synthetic OHCA hotspots per Sprint 26; 37 SA2 centroids, Greater Sydney. Station locations are observed; as
Projection: equirectangular, scaled at the latitude midpoint. Base-map opacity 0.75 (Sprint 35). Coverage rings omitted from this figure; the dist
Figure by Grant (mapping agent), LIFT Study 01 · Sprint 40 · 2026-04-22. Pending operational-siting validation against NSW Ambulance statio

Figure 3.3. NSW Ambulance stations inside the Greater Sydney bounding box as candidate drone-base locations, plotted over the SEIFA-weighted synthetic OHCA hotspot surface. Top-10 stations selected by greedy set-cover at 6 km service radius are filled in signal-red and labelled. Data label: observed (station locations) + synthetic (demand surface).

Running the same greedy set-cover algorithm over the three candidate pools, at the 6 km base-scenario service radius, produces the comparison in Table 3.1.

Table 3.1. Greater Sydney OHCA coverage by candidate pool, 6 km service radius, greedy set-cover, 10-base fleet. Denominator is the 4,800-event annual city envelope. Data label: synthetic demand + observed station locations. {#tbl:candidate-pool-comparison}

Candidate pool	Bases needed to saturate	Peak coverage (of 4,800)
SA2 synthetic centroids (baseline)	6	50.9% [ESTIMATED]
NSW Ambulance stations	10	69.7% [ESTIMATED]
Hybrid (either pool)	10	69.7% [ESTIMATED]

The full base-by-base rollout for each pool and each radius scenario lives in the coverage-results dataset (reproducibility appendix) .

Key insight. Anchoring candidate bases on the existing ambulance-station network lifts peak coverage from roughly 51% to roughly 70% at the 6 km base radius, because the station network is denser and more geographically spread than the 10-hotspot synthetic pool, so the greedy rollout has more productive candidates to choose from before saturation sets in. The hybrid pool does not beat the station-only pool at 6 km, because the first ten hybrid picks are all stations; at the optimistic 8 km radius the hybrid pool edges past station-only (85% versus 79% of city envelope) because the synthetic centroids fill a small Campbelltown-south gap the station network does not cover as tightly. The result should be read as a sensitivity pass, not a re-baselining: the synthetic-centroid numbers remain the headline planning envelope because they are the pool the project has calibrated against, and the station-pool numbers depend on NSW Ambulance agreeing to co-location, which is a negotiation rather than a technical fact.

Caveats and next steps

Both maps draw on the synthetic cohort only; neither represents observed NSW Ambulance incident data. Under the current methodology the synthetic proxy is SEIFA-IRSD-weighted; the derivation, the Hasselqvist-Ax 2015 base-rate anchor, and the pending-verification labelling convention are documented in Appendix M: Methodology: synthetic OHCA proxy with SEIFA IRSD weighting. The synthetic counts remain above the n < 5 suppression threshold for every SA2, so no cells are suppressed, but rates are not age-standardised and will need Empirical Bayes smoothing when the real cohort is slotted in. The nearest-centroid raster proxy used in Figure 3.1 is a presentational stand-in for ABS SA2 polygons and is explicitly labelled as such on the map itself. The final report build will swap in the true SA2 geometry and re-run the quantile classification; the class breaks will shift slightly but the south-west cluster is expected to survive the substitution because the ranks are carried by the underlying rates, not by the polygon shapes.

CHAPTER 04

The Orange Paradox

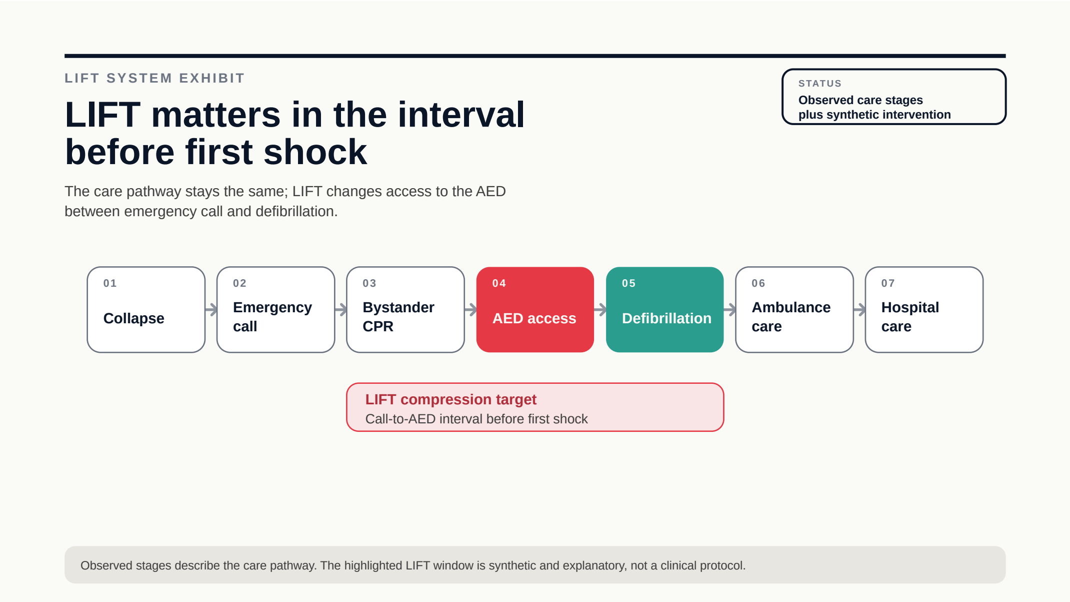

LIFT SYSTEM EXHIBIT

LIFT matters in the interval before first shock

The care pathway stays the same; LIFT changes access to the AED between emergency call and defibrillation.

STATUS
Observed care stages plus synthetic intervention

01 **Collapse** → 02 **Emergency call** → 03 **Bystander CPR** → 04 **AED access** → 05 **Defibrillation** → 06 **Ambulance care** → 07 **Hospital care**

LIFT compression target
Call-to-AED interval before first shock

Observed stages describe the care pathway. The highlighted LIFT window is synthetic and explanatory, not a clinical protocol.

Out-of-hospital cardiac arrest kills roughly 30,000 Australians each year [3]. That is more than breast cancer. More than prostate cancer. More than road trauma. In Greater Sydney alone, the model estimates approximately 4,800 events per year [ESTIMATED], concentrated not in public spaces but inside private homes. Over 70% of OHCA occurs in residential settings [3].

> *The majority of out-of-hospital cardiac arrests in Australia occur in private residences, where publicly accessible AEDs are rarely available.*
>
> — AIHW, *OHCA Annual Report* [3]

Over 70%

Of attended OHCA events in NSW occur in private residences where public AEDs are not accessible. Fixed AED placement strategies do not reach these patients [3,4].

The clinical science is settled. For every minute without defibrillation, survival from witnessed ventricular fibrillation drops 7-10% [5,6]. The American Heart Association recommends defibrillation within 3-5 minutes of collapse [1]. Bystander defibrillation before EMS arrival is associated with 53% 30-day survival for shockable rhythms. Without it: 16% [7]. That is the difference between a patient who walks out of hospital and a patient who does not leave.

And yet.

In at least one regional NSW community, the aspiration has been to place an AED within 200 metres of every resident. In Orange, the Heart of the Nation charity, founded by cardiac-arrest survivor Greg Page after his 2020 collapse, worked with Orange City Council to install defibrillators in every town park, with councillors unanimously backing the plan [8,9]. The ambition is sound: saturate the public AED footprint so that no cardiac arrest occurs beyond arm's reach of the one device proven to change outcomes. On its face, this is the obvious solution. Put enough devices in enough places and the problem disappears.

It does not disappear. It reveals a paradox with three dimensions.

Scale and cost

What is achievable in a compact regional town does not transfer to a metropolitan area of five million people spread across 12,000 square kilometres. Placing an AED within 200 metres of every resident in a regional centre requires a manageable number of devices in a concentrated geography. Replicating that density across Greater Sydney, from Penrith to Cronulla, from Campbelltown to the Northern Beaches, would demand tens of thousands of units, each requiring purchase, installation, signage, registration, and ongoing maintenance including battery replacement and pad expiry [ESTIMATED]. The capital outlay alone would run into the tens of millions. Maintaining currency across that fleet, ensuring every device is charged, in-date, and functional at the moment it is needed, would constitute a programme of work in its own right.

Scaling to a metropolitan area is qualitatively different from scaling within a compact regional town; the coordination cost grows non-linearly, the maintenance burden compounds, and the marginal return diminishes. The last 10% of coverage, covering cul-de-sacs, acreage blocks, and isolated pockets between commercial strips, costs disproportionately more than the first 90%.

Awareness, access, and use

Even where AEDs exist, three barriers stand between the device and the patient.

Awareness. The bystander who witnesses a cardiac arrest must know that an AED is nearby. Most do not. Public-access AED registries exist, but cardiac arrest does not wait for someone to consult a database. The 000 call-taker may direct a bystander to the nearest registered device, but the bystander is already managing CPR, managing panic, managing a dying person on the floor. Knowledge of AED locations is not widely held in the community, and under the cognitive load of a cardiac arrest, even known information becomes difficult to retrieve.

Access. The AED must be physically reachable. An AED mounted inside a surf club is not reachable at 06:42 on a Tuesday morning. An AED in a school is not reachable before 08:30 or after 15:30 or during holidays. An AED in a GP surgery is not reachable on weekends. An AED in a community centre 400 metres away requires a round trip of five to eight minutes on foot, and that assumes the bystander leaves the patient to go looking [OBSERVED]. Five AEDs within 1.5 kilometres, none accessible. That is the pattern described in the preface of this report, replicated across suburban Sydney every week.

Willingness and competence to use. Even when a bystander locates and reaches an AED, they must be willing to apply it to a human being in extremis. Public anxiety about using AEDs remains high despite their design for lay use [5]. Fear of causing harm, uncertainty about legal liability, and unfamiliarity with pad placement are all documented barriers to bystander defibrillation in the literature. The device is forgiving; the psychology is not.

Measurable outcomes

The most difficult dimension. Does saturating an area with AEDs produce measurable improvements in survival? Proximity alone does not complete the chain of survival. A defibrillator that is purchased, installed, registered, maintained, located by a bystander, accessed within minutes, and correctly applied can save a life. Remove any link and the chain fails. AEDs work, unambiguously [7]; the open question is whether a fixed-installation strategy translates proximity into actual use, and actual use into population-level survival gains.

The evidence is mixed. Public-access defibrillation programmes in high-traffic locations such as airports, casinos, and sports arenas have demonstrated clear survival benefit [5]. These are environments with trained staff, high footfall, and rapid recognition. Residential settings share none of these characteristics. The patient collapses at home, often in the early morning, often alone or with a single household member who may themselves be elderly. The nearest AED is behind a locked door. The chain of survival breaks at the same link every time: the device does not reach the patient.

The fire extinguisher standard

Consider the fire extinguisher. Building codes require extinguishers within 30 metres of travel distance in most occupied spaces. Not because fires are common, but because when they happen, seconds determine whether a person walks out or does not. The reasoning is simple: if the device exists but cannot be reached in time, it may as well not exist.

Cardiac arrest is the residential equivalent. The patient has minutes, the intervention exists, and the gap is in delivery rather than knowledge or technology.

The paradox, then, is this: the aspiration to place an AED within 200 metres of every resident is clinically sound but operationally insufficient. It cannot scale to a metropolitan area at reasonable cost. Where it does exist, awareness, access, and willingness erode the benefit. And even where every barrier is cleared, the evidence for population-level survival gains from fixed residential AED installations remains inconclusive.

This is a system design problem. The devices exist, the evidence for early defibrillation is unambiguous, and the clinical protocols are established. What is missing is a delivery mechanism that matches the geography and timing of the events themselves: residential, dispersed, and time-critical. The following section examines how current emergency response systems perform against that requirement and where they fall short.

CHAPTER 05

06:42: A Tuesday Morning in Campbelltown

Campbelltown South is an SA2 in Sydney’s outer southwest. Disadvantage decile 1. Population 11,200. Detached housing: 82%. Residents aged 65 and over: 14.5% [10]. Streets of single-storey brick and tile homes set back from the kerb. Driveways with two cars. Low front fences. Quiet before seven.

06:42. A 71-year-old man steps out of the bedroom and collapses in the hallway [SYNTHETIC]. His wife hears the fall from the kitchen. She finds him unresponsive on the carpet, not breathing normally. She does not know what cardiac arrest looks like. She knows something is wrong.

06:43. She calls 000. The dispatcher recognises the pattern within 40 seconds: unresponsive, abnormal breathing. The call is coded as suspected cardiac arrest. Two things happen simultaneously. The dispatcher begins telephone CPR instructions, talking the wife through chest compressions on the hallway floor. And the dispatch system triggers a drone from the nearest candidate base, approximately 5 kilometres away [SYNTHETIC].

06:44. The drone is airborne. Cruise speed: 100 km/h [ESTIMATED]. It flies a direct path, unimpeded by traffic lights, school zones, or the roundabout on Narellan Road. The wife is on her knees doing compressions. They are shallow and fast. The dispatcher corrects her tempo.

06:47. The drone arrives overhead. Audio and visual alerts activate. The AED descends to the front yard. Three minutes from launch.

06:48. The wife retrieves the AED. She opens it. The device talks to her. Peel the pads. Place them on his bare chest. Analysing. Shock advised. She presses the button. First shock delivered at approximately six minutes from collapse [ESTIMATED].

At six minutes, the survival probability for a shockable rhythm is roughly 36% [ESTIMATED, derived from [5]; $P = 0.67 \times \exp(-0.10 \times t)$, $t = 6$].

06:53. The ambulance arrives. Paramedics take over with advanced life support. The AED has recorded two shocks and six minutes of CPR data. The patient has a pulse.

Now remove the drone.

06:42. The same collapse. The same hallway.

06:43. The same 000 call. The same telephone CPR. No drone dispatch.

06:53. The ambulance arrives carrying the first defibrillator. Eleven minutes from collapse. The wife has been doing compressions for ten minutes. She is exhausted. Compression depth has dropped.

At eleven minutes, survival probability is roughly 22% [ESTIMATED, same model, t = 11].

The difference: 14 percentage points. In a population where roughly one in five attended OHCA presents with a shockable rhythm [3], that difference is not abstract.

The drone did not replace the ambulance. The paramedics still arrived at 06:53. The clinical handover still happened. What changed was the six minutes between collapse and first shock. Those six minutes are the entire intervention.

This scenario is synthetic. The SA2 is real. The survival curve is evidence-based. The question it raises is not.

CHAPTER 05A

GoodSAM and the trained-responder layer

The preface to this report is a first-person account of two GoodSAM activations in Sydney, both without a reachable defibrillator in time. That chapter is the lived experience. This one is the programme evidence that sits underneath it. Trained-bystander-responder systems are already operational in NSW and Victoria. They work. They also reveal, with some precision, the gap that an AED-delivery layer is meant to close.

What GoodSAM is

GoodSAM (Good Smartphone Activated Medics) is a smartphone application that notifies registered off-duty responders (paramedics, nurses, doctors, CPR-trained laypeople) when a suspected cardiac arrest is reported to the emergency-call service within a defined radius. The responder can accept, decline, or ignore the alert. On acceptance, the app routes them to the patient and, where relevant, to the nearest registered public-access AED. Dispatch of the ambulance is parallel, not contingent on the alert. The app originated in the United Kingdom and has since been adopted by several Australian state ambulance services, as well as services in Europe and North America.

Where it is deployed

In Victoria, Ambulance Victoria launched GoodSAM in 2018 alongside a public AED registry [11]. In New South Wales, NSW Ambulance launched the programme on 23 November 2023; at the 12-month review it had 8,372 registered responders with 1,745 recorded on-scene arrivals [12]. By early 2026 the NSW roll had grown past 13,500 registered responders and the service had publicly credited the programme with 100 lives saved [13]. Comparable systems run in London, the English East Midlands, Stockholm, Copenhagen, and parts of the Netherlands and Germany.

What the evidence says about clinical impact

Three peer-reviewed signals are worth citing directly. A Stockholm randomised trial of mobile-phone dispatch found that bystander-initiated CPR rose from 48% in the control arm to 62% in the intervention arm, a 14-point absolute increase from a single software layer [14]. A 2022 analysis of GoodSAM across London and the English East Midlands (5,237 confirmed OHCAs) reported that when a responder accepted an alert, the adjusted odds of survival to hospital discharge were 3.15× in London (95% CI 1.19–8.36) and 3.19× in the East Midlands (95% CI 1.17–8.73) [15]. The most recent Victorian evidence, a 2018–23 cohort of 9,196 OHCAs, found that when a smartphone-activated volunteer arrived before EMS, the patient was 8× more likely to receive bystander CPR, 16× more likely to receive bystander defibrillation, and had 37% higher odds of survival to hospital discharge [16]. Across the same two decades in Victoria, the share of OHCAs receiving bystander CPR rose from 40.3% in 2003–04 to 72.2% in 2021–22; Utstein-comparator survival to discharge roughly tripled [11]. Earlier Victorian work on the AED-registry layer sits alongside these findings as a precedent for public-access defibrillation in the same jurisdiction [17].

> *When a responder accepted an alert, the adjusted odds ratio for survival to hospital discharge was 3.15 in London and 3.19 in the East Midlands.*
>
> — Smith et al., *European Heart Journal — Acute Cardiovascular Care* [15]

13,500 responders

Registered on the NSW Ambulance GoodSAM platform by early 2026, with the programme publicly credited with 100 lives saved since its November 2023 launch [12,13].

The Sydney gap this paper addresses

Two points matter for the rest of the argument. The first: the responder layer is built. A trained person can be in the hallway inside three to five minutes of an alert in a well-populated part of Sydney, as Chapter 01 records. The second: what that responder cannot do, in more than seven out of ten Sydney residential arrests, is arrive with a defibrillator in their hand. Public-access AEDs cluster in the places where GoodSAM responders are *least* likely to need them (commercial strips, transport interchanges, sports venues), and are behind locked doors at night and in the early morning, which is when a large share of residential events occur [3]. The outcome gap is payload delivery, not awareness or willingness.

How drone-AED delivery complements, rather than replaces

GoodSAM gets the responder to the scene. A drone-AED layer gets the defibrillator to the scene. The two vectors run in parallel and arrive independently; neither one replaces the ambulance, and neither one replaces the other. International evidence has already shown that airborne AED delivery can beat EMS to the patient in more than two-thirds of suspected OHCAs in a live operational setting [18], and at least one documented save now exists in which the drone-delivered device was used by a bystander before EMS arrival [19]. The question this paper poses is narrower and more local: given that Sydney already has a trained-responder layer and an AED registry, what would it take to add a dispatched-defibrillator layer, and what would the incremental survival benefit be in the residential geography where GoodSAM is already active?

That is the question the chapters that follow address.

CHAPTER 06

Demand Signal

Out-of-hospital cardiac arrest does not strike uniformly. It clusters, driven by age, population size, and social disadvantage. To estimate where burden falls across Greater Sydney, we built a synthetic demand surface across 37 SA2 areas, selected for geographic spread and demographic diversity [3,10].

NSW records approximately 4,800 OHCA events per year in Greater Sydney [ESTIMATED from state-level reporting]. No publicly available dataset geocodes these events to suburb level. We allocated an annual total of roughly 4,200 cases across the 37 project SA2s using the SEIFA-weighted proxy documented in the Methodology chapter: event counts are proportional to population, to the age-65+ share, and to a SEIFA IRSD decile weight that is highest in the most disadvantaged areas [10,20]. The resulting event counts are synthetic; they reflect demographic plausibility under an explicit formula, not verified incident locations [SYNTHETIC].

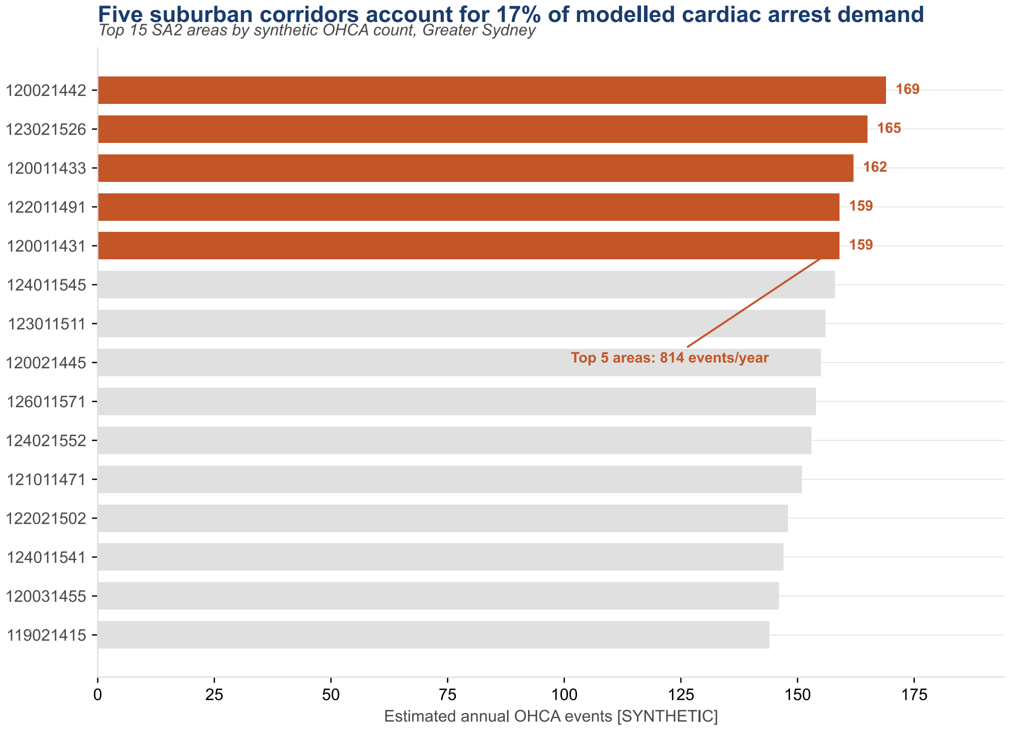

Data label: synthetic/estimated. Source: synthetic_ohca_hotspots.csv

2 Figure 6.1 — Synthetic OHCA counts are concentrated in the highest-ranked proxy hotspot areas. Each bar is one SA2 in the Greater Sydney 37-area proxy cohort, ordered by SEIFA-weighted synthetic count. Data label: synthetic.

OHCA rates concentrate in Sydney's disadvantaged south-west once SEIFA IRSD is weighted in

Synthetic OHCA incidence per 100,000 population · 5-class quantile · 37 SA2s, Greater Sydney · proxy pending NSW Ambulance data

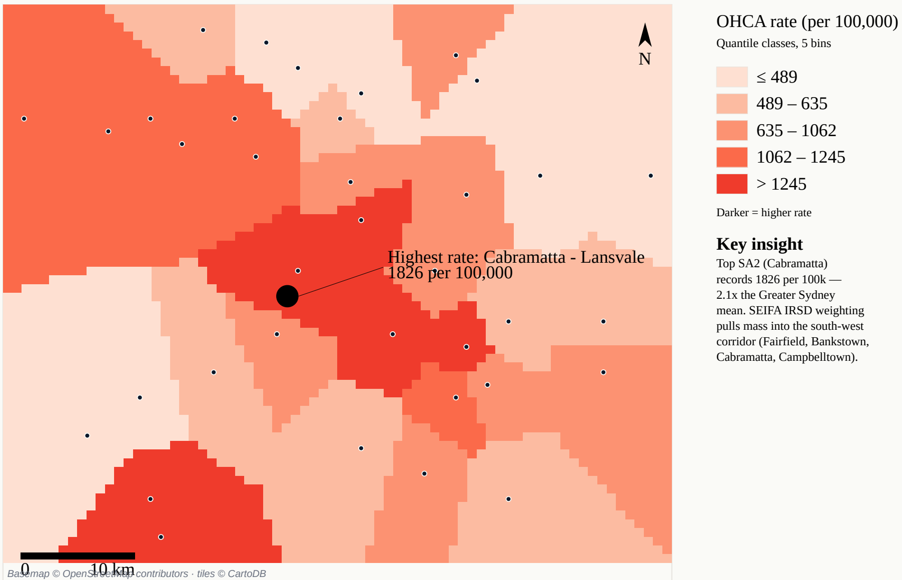

Key insight

Top SA2 (Cabramatta) records 1826 per 100k — 2.1x the Greater Sydney mean. SEIFA IRSD weighting pulls mass into the south-west corridor (Fairfield, Bankstown, Cabramatta, Campbelltown).

Source: synthetic OHCA hotspots (data/derived/synthetic_ohca_hotspots.csv), derived from ABS SA2 2021 geography (data/curated/sa2_greate
Areal unit: ABS Statistical Area Level 2 (SA2), Greater Sydney (n=37). MAUP caveat: spatial pattern depends on the choice of SA2 boundaries
Projection: equirectangular, scaled at the latitude midpoint. Choropleth surface: nearest-centroid (Voronoi) raster proxy rendered at 72x64 cells
Small-numbers caveat: all SA2s carry ≥ 33 synthetic cases, above the suppression threshold; Empirical Bayes smoothing remains advisable at fi
Figure by Grant (mapping agent), LIFT Study 01 · Sprint 20 · re-rendered Sprint 26 (SEIFA IRSD weighting) · 2026-04-21.

3 Figure 6.2 — SA2 demand allocation score (SEIFA-IRSD-weighted), ABS SA2 geography, 5-class quantile. Under SEIFA IRSD weighting the darkest band sits across Sydney’s south-west corridor (Fairfield, Cabramatta, Bankstown, Campbelltown). The surface is a modelled allocation score for base-placement prioritisation, not an observed incidence rate; see Appendix M for the allocation-vs-incidence distinction. Data label: modelled.

Fairfield / Bankstown / Cabramatta top the SEIFA-weighted synthetic OHCA proxy

Proportional circle (area ∝ annual cases) · 37 SA2 centroids · Greater Sydney synthetic cohort · proxy pending NSW Ambulance data

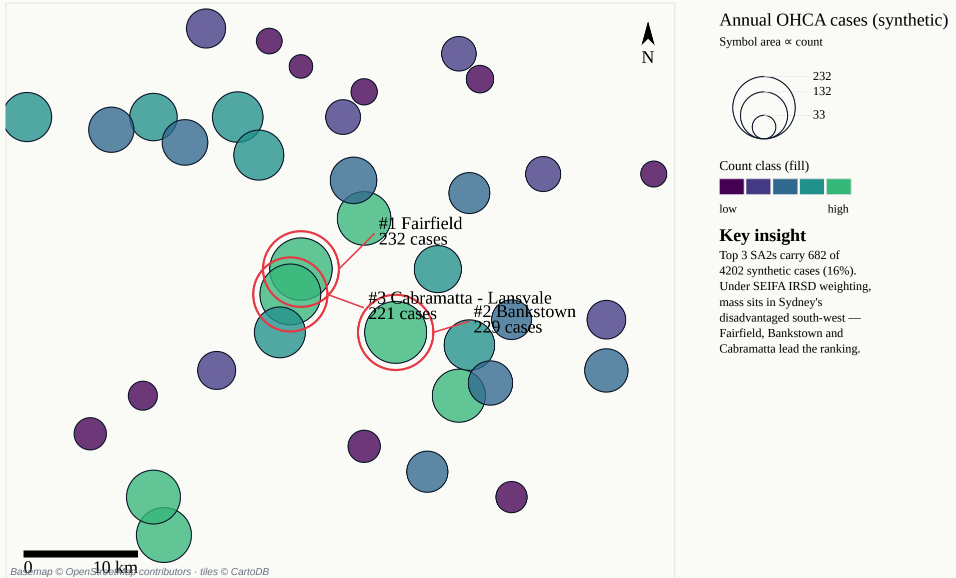

Key insight

Top 3 SA2s carry 682 of 4202 synthetic cases (16%). Under SEIFA IRSD weighting, mass sits in Sydney's disadvantaged south-west — Fairfield, Bankstown and Cabramatta lead the ranking.

Source: data/derived/synthetic_ohca_hotspots.csv (synthetic annual counts) · SA2 centroids from data/curated/sa2_greater_sydney.csv (ABS 20
Areal unit: ABS Statistical Area Level 2 (SA2) centroid (n=37), Greater Sydney. MAUP note: counts aggregate all within-SA2 events to a singl
Projection: equirectangular, scaled at the latitude midpoint. Symbol area is linearly proportional to count (no Flannery perceptual correction app
Small-numbers caveat: smallest symbol represents 33 synthetic cases; all symbols exceed the n<5 suppression threshold.
Figure by Grant (mapping agent), LIFT Study 01 · Sprint 20 · re-rendered Sprint 26 (SEIFA IRSD weighting) · 2026-04-21.

4 Figure 6.3 — Proportional-symbol map of synthetic annual OHCA counts by SA2 centroid, Greater Sydney (n = 37). Symbol area is linearly proportional to count; the three largest circles sit in the south-west corridor. Data label: synthetic.

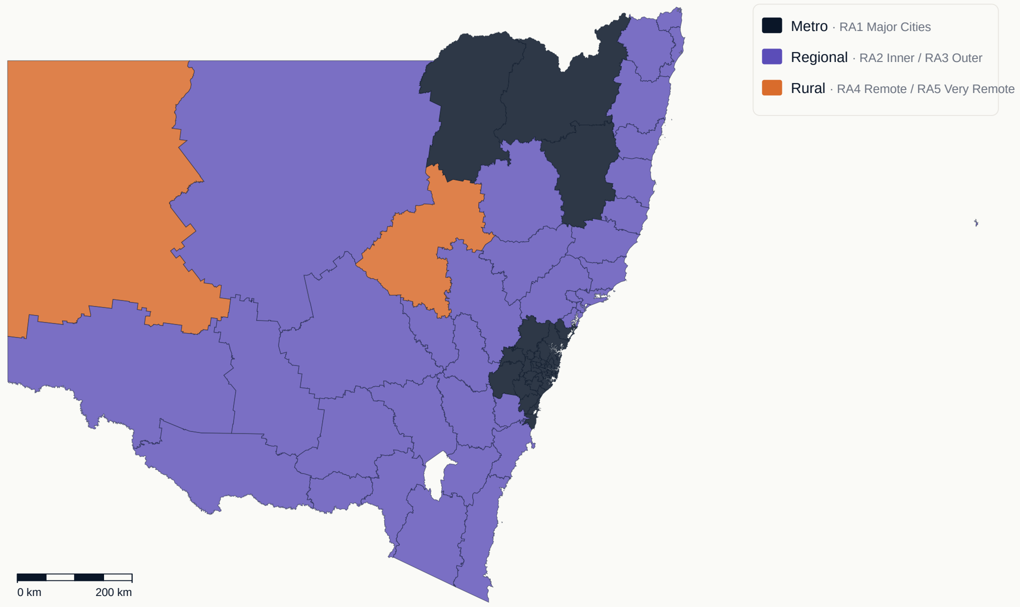

NSW SA3 three-tier Remoteness Area choropleth

Each polygon is one ABS SA3, filled by predominant Remoteness Area tier (Metro / Regional / Rural).

Aggregation: SA3 (ABS ASGS 2021) · Output status: derived

Source: ABS ASGS Edition 3 SA3 2021 + RA2021 (Remoteness Area). © Commonwealth of Australia (CC BY 4.0).

5 Figure 6.4 — NSW SA3 three-tier polygon choropleth (Metro / Regional / Rural). Each polygon is one ABS SA3 filled by its predominant ABS Remoteness Area tier (RA1 = Metro; RA2/RA3 = Regional; RA4/RA5 = Rural). The southwestern Sydney corridor sits in Metro; the Far West NSW SA3 sits in Rural; mid-coast and inland SA3s show as Regional. Aggregation: SA3 (ABS ASGS 2021). Data label: derived.

The five highest-burden SA2 areas under this model:

Table 6.1 — Estimated annual OHCA events by SA2 area, top five under the SEIFA-weighted proxy (source: Appendix M). Data label: synthetic. {#tbl:top5-hotspots}

SA2 Area	Allocated Annual Events
Fairfield	232 [SYNTHETIC]
Bankstown	229 [SYNTHETIC]
Cabramatta – Lansvale	221 [SYNTHETIC]
Campbelltown – South	180 [SYNTHETIC]
Campbelltown – North	173 [SYNTHETIC]

A pattern emerges. The five SA2s carrying the highest estimated burden all sit in IRSD decile 1 or 2 (the most disadvantaged end of the SEIFA distribution), and each carries a sizeable resident population with a meaningful age-65+ share [10,20]. Disadvantage correlates with higher rates of chronic cardiovascular disease and, in many areas, with longer ambulance response times. Older populations carry elevated cardiac arrest risk independent of geography.

These top five SA2s alone account for roughly 1,035 of the ~4,200 allocated annual events, close to a quarter of the total demand concentrated in about an eighth of the study areas [SYNTHETIC].

A transparency note on what this surface can and cannot show. The proxy is a model, not a measurement. It is built from ABS 2021 Census inputs and a linear SEIFA IRSD decile weighting; it is *not* de-identified NSW Ambulance incident data, and every absolute count on the page should be read pending that validation. What a proxy like this *can* show is the demographic shape of risk: which SA2s concentrate the joint load of population, age, and disadvantage the published literature associates with OHCA [3]. What it *cannot* yet show is the actual spatial point pattern of incidents, the street-level timing, or whether a protective factor (preventive cardiology access, earlier bystander response) partially offsets demographic risk in a given suburb. Resolving the gap between the two is a validation task against NSW Ambulance records, flagged in the Methodology chapter and in chapter 11’s risk register; relative SA2 ordering is expected to hold under replacement with the real cohort, absolute counts are not.

This demand surface is not dispatch-grade. It cannot tell us which street will see the next cardiac arrest. What it can do is identify the zones where drone positioning would intercept the greatest number of events, given demographic realities. For feasibility scoping (deciding where to place a small number of bases and estimating population-level coverage) this resolution is sufficient. Operational deployment would require geocoded incident data from NSW Ambulance, a step beyond the boundary of this analysis.

The surface is a planning tool. It tells us roughly where to look, not precisely where to land.

CHAPTER 07

Coverage and Staged Rollout

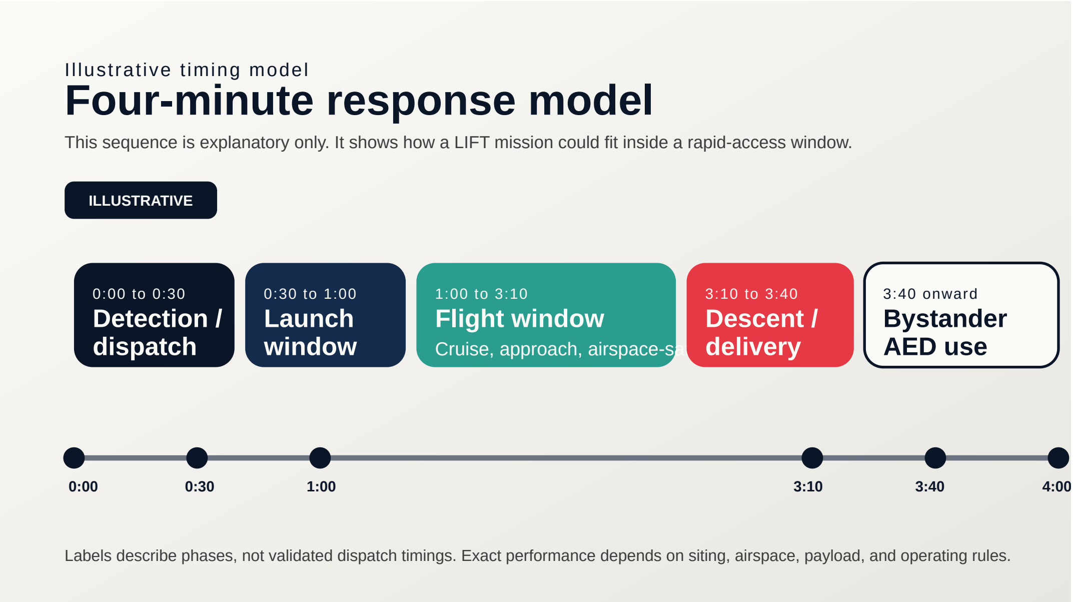

Illustrative timing model

Four-minute response model

This sequence is explanatory only. It shows how a LIFT mission could fit inside a rapid-access window.

ILLUSTRATIVE

0:00 to 0:30 **Detection / dispatch**

0:30 to 1:00 **Launch window**

1:00 to 3:10 **Flight window** Cruise, approach, airspace-s

3:10 to 3:40 **Descent / delivery**

3:40 onward **Bystander AED use**

Labels describe phases, not validated dispatch timings. Exact performance depends on siting, airspace, payload, and operating rules.

A drone on a launchpad covers nothing. Coverage depends on where bases are placed and how far each drone can fly within the critical response window. We modelled base placement using a greedy set-cover algorithm: place the first base where it captures the most OHCA events, then place each subsequent base to maximise the incremental gain [ESTIMATED].

Using a 6 km service radius as the base scenario:

Table 7.1 — Cumulative OHCA coverage by number of drone bases, 6 km service radius. Data label: estimated. {#tbl:coverage-by-bases}

Bases	Cumulative OHCA Coverage
1	12.8% [ESTIMATED]
2	21.3% [ESTIMATED]
3	28.0% [ESTIMATED]
5	40.0% [ESTIMATED]
10	51.4% [ESTIMATED]

The full coverage-by-base series across conservative / base / optimistic scenarios is carried in the coverage-results dataset (reproducibility appendix).

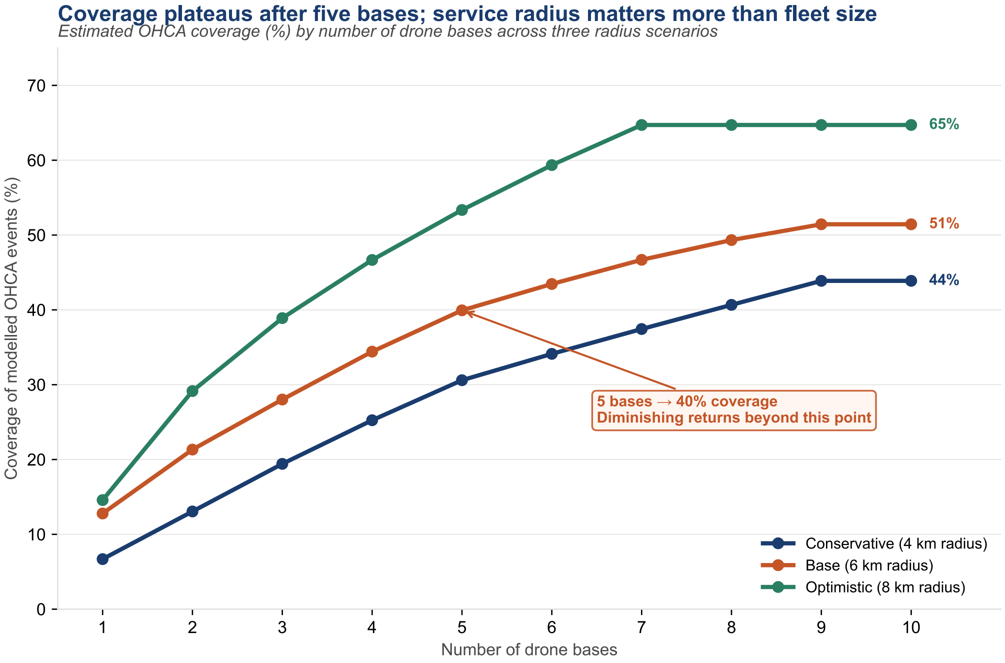

Data label: estimated. Source: coverage_vs_base_count_table.csv

6 Figure 7.1 — Cumulative coverage rises as additional candidate bases are added in the synthetic feasibility model, flattening after the fifth base in every radius scenario. Data label: synthetic/estimated.

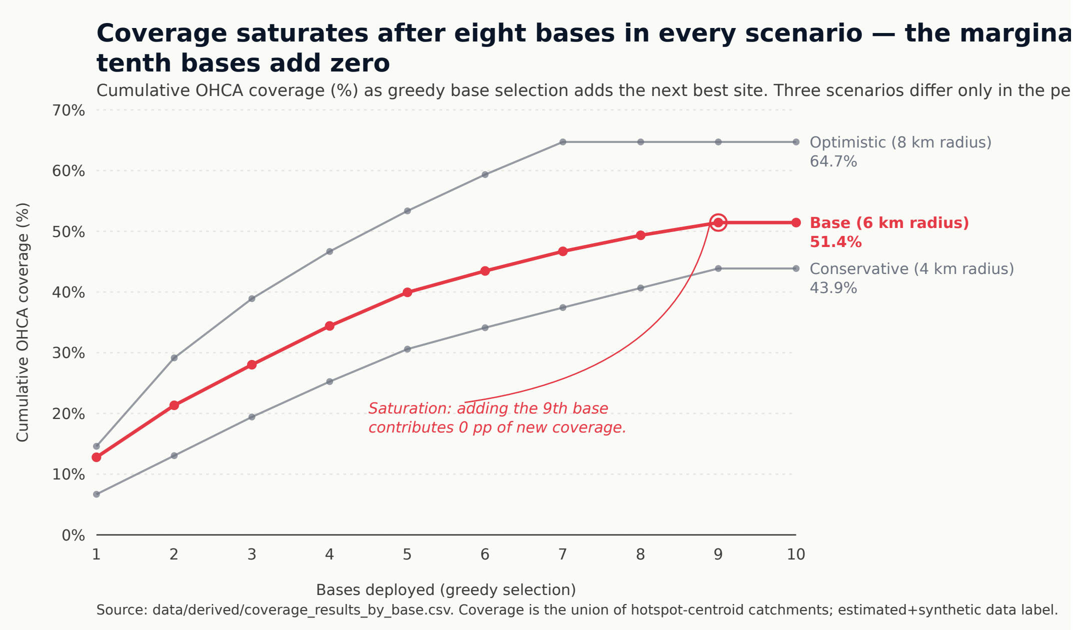

Coverage saturates after eight bases in every scenario — the margina
tenth bases add zero

Cumulative OHCA coverage (%) as greedy base selection adds the next best site. Three scenarios differ only in the pe

Source: data/derived/coverage_results_by_base.csv. Coverage is the union of hotspot-centroid catchments; estimated+synthetic data label.

7 Figure 7.1b — Cumulative OHCA coverage vs number of drone bases across three radius scenarios (4 km conservative, 6 km base, 8 km optimistic). Marginal coverage collapses beyond the eighth base. Source: the coverage-results dataset (reproducibility appendix). Data label: synthetic/estimated.

The first two bases deliver 21.3% coverage. Doubling to five bases does not double coverage; it reaches 40.0%. By ten bases, each additional site adds only two to three percentage points. Diminishing returns set in after five.

We tested three radius scenarios. Under a conservative 4 km radius, reflecting headwinds, airspace restrictions, or heavier payloads, coverage contracts at every fleet size. Under an optimistic 8 km radius, coverage expands, but the shape of the diminishing-returns curve stays the same. Across all scenarios, service radius is the dominant parameter. Whether a drone can reach 6 km or only 4 km matters more than whether the network has seven bases or ten. This finding directs attention toward aircraft performance and airspace access as the binding constraints on real-world effectiveness.

All radii are geometric great-circle distances from base to SA2 centroid. They do not account for controlled airspace boundaries, terrain obstacles, weather, or flight-path routing. Operational coverage will be smaller than these estimates.

Coverage overlay: 10-base footprint over the SEIFA-weighted OHCA surface

Candidate bases anchored on Sprint 26 top-10 hotspots · 4 km and 5 km service radii shown · 37 SA2 centroids · Greater Sydney synthetic coho

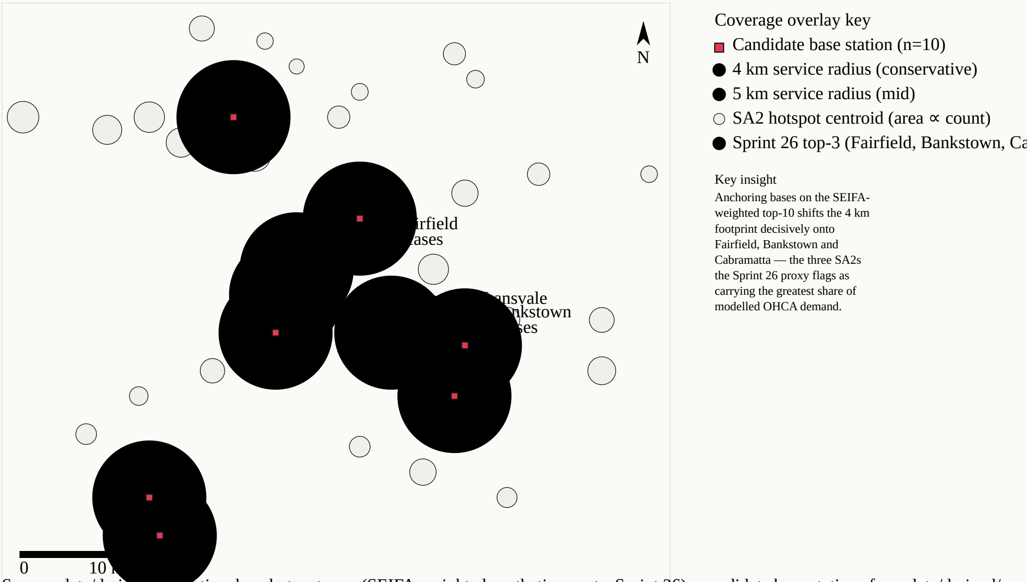

Key insight

Anchoring bases on the SEIFA-weighted top-10 shifts the 4 km footprint decisively onto Fairfield, Bankstown and Cabramatta — the three SA2s the Sprint 26 proxy flags as carrying the greatest share of modelled OHCA demand.

Source: data/derived/synthetic_ohca_hotspots.csv (SEIFA-weighted synthetic counts, Sprint 26) · candidate base stations from data/derived/can
Areal unit: ABS SA2 centroids (n=37), Greater Sydney. Base markers are inferred from hotspot centroids rather than operational siting data; dat
Projection: equirectangular, scaled at the latitude midpoint. Coverage circles are great-circle radii (4 km solid, 5 km dashed); they do NOT accou
Coverage footprint covers the Sprint 26 top-3 (Fairfield, Bankstown, Cabramatta–Lansvale) inside the 4 km conservative ring when bases are si
Figure by Grant (mapping agent), LIFT Study 01 · Sprint 20 · regenerated Sprint 30 (SEIFA-anchored base stations) · 2026-04-21.

8 Figure 7.2 — Coverage overlay: candidate base stations and 4 km / 5 km service circles plotted over the SEIFA-weighted synthetic demand surface. Top-3 hotspots (Fairfield, Bankstown, Cabramatta) sit inside the base-scenario coverage footprint. Data label: synthetic/inferred.

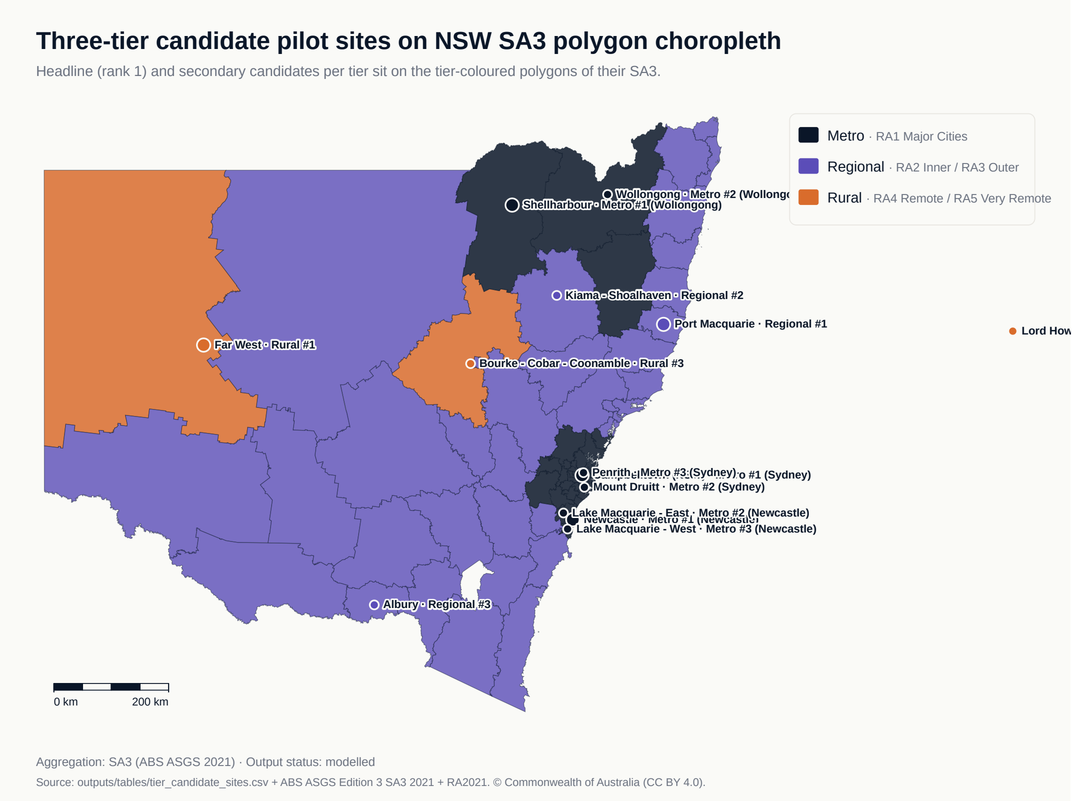

Three-tier candidate pilot sites on NSW SA3 polygon choropleth

Headline (rank 1) and secondary candidates per tier sit on the tier-coloured polygons of their SA3.

Aggregation: SA3 (ABS ASGS 2021) · Output status: modelled

Source: outputs/tables/tier_candidate_sites.csv + ABS ASGS Edition 3 SA3 2021 + RA2021. © Commonwealth of Australia (CC BY 4.0).

9 Figure 7.3 — Three-tier candidate pilot sites overlaid on the NSW SA3 tier polygon choropleth. Headline candidates (rank 1) per tier sit on the corresponding tier-coloured polygons; secondary candidates fill out the picks. Aggregation: SA3 (ABS ASGS 2021). Data label: modelled.

Staged rollout

The coverage model supports a staged deployment, not because caution is a virtue in itself, but because sequencing is the fastest way to generate usable evidence.

> The Sydney Harbour Bridge was not built in one span. It was constructed from both sides simultaneously, with each section tested before the next was added. A drone AED network follows the same logic: start with two bases in the highest-need areas, validate the approach, then extend.

Stage 1 positions two bases in the highest-burden SA2 clusters. Two bases cover an estimated 21.3% of annual events [ESTIMATED], enough to generate statistically meaningful response-time data within 12 to 18 months. This stage validates flight operations, dispatch integration, and real-world delivery times against the geometric estimates.

Stage 2 extends to five bases, guided by Stage 1 operational data. If actual service radii prove shorter than 6 km, the algorithm re-optimises placement. If certain areas show access constraints (controlled airspace near Kingsford Smith, for instance), the model reallocates those bases. Five bases at validated parameters are worth more than ten placed on unchecked assumptions.

Stage 3 scales toward ten bases only if Stage 2 demonstrates worthwhile incremental gains at confirmed operational radii.

This is risk reduction through sequencing. Each stage produces data that sharpens the next.

If bases are co-located at ambulance stations

The coverage numbers above use a synthetic pool of candidate base locations anchored on the top-10 SEIFA-weighted OHCA hotspots (chapter 03). A more operationally realistic candidate pool is the existing NSW Ambulance station network: co-locating drones with paramedic crews reuses electrical supply, 24/7 security, and site access, and it simplifies dispatch integration. Chapter 03 describes the station pool; the base-by-base rollout lives in the coverage-results dataset (reproducibility appendix).

Running the same greedy set-cover algorithm against the 44 NSW Ambulance stations inside the Greater Sydney bounding box at 6 km service radius, the staged-rollout picks are:

Table 7.2. Ambulance-station candidate picks under 6 km greedy set-cover. Cumulative coverage is against the 4,800-event annual city envelope. Data label: synthetic demand + observed station locations. {#tbl:station-staged-rollout}

Stage	Station picks	Cumulative city coverage
Stage 1 (3 stations)	Cabramatta, Rockdale, Mount Druitt	31.6% [ESTIMATED]
Stage 2 (6 stations)	adds Greenacre, Campbelltown, Blacktown	52.8% [ESTIMATED]
Stage 3 (10 stations)	adds Parramatta, Randwick, Sutherland, Castle Hill	69.7% [ESTIMATED]

The staged progression above covers a similar geography to the synthetic-centroid plan, with south-western Sydney dominating the first three picks, but routes every base through an existing operational site rather than a greenfield cradle. That distinction matters for the cost envelope (chapter 09), the operating-model integration (chapter 10), and the risk register (chapter 11). It does not displace the synthetic-centroid numbers used as the planning headline; it sits alongside them as a realistic deployment variant pending NSW Ambulance agreement to co-location.

CHAPTER 08

Survival Impact

Cardiac arrest survival is a race measured in minutes. The relationship is not linear but exponential decay.

The model follows the form $P = 0.67 \times \exp(-0.10 \times t)$, where P is survival probability and t is minutes from collapse to first defibrillation. This draws on Valenzuela et al. 1997 [5] and Larsen et al. 1993 [21], both establishing the steep, time-dependent decline in survival for ventricular fibrillation and pulseless ventricular tachycardia. It is one of the most replicated findings in resuscitation science.

The numbers at specific response times:

At 4 minutes, the modelled drone delivery target, survival probability sits in the **40–50%** range once the ±20% decay-parameter sensitivity is applied. At 8.5 minutes, the Greater Sydney Cat-1 ambulance response median, it falls to the **24–34%** range. At the Rural-tier Cat-1 response median of roughly 18 minutes, it drops below 11% — the survival window has effectively closed before ambulance arrival, which is the substantive reason the rural-tier case is so much more leveraged than metro. Every additional minute costs roughly 7–10 percentage points. The curve does not negotiate. [ESTIMATED]

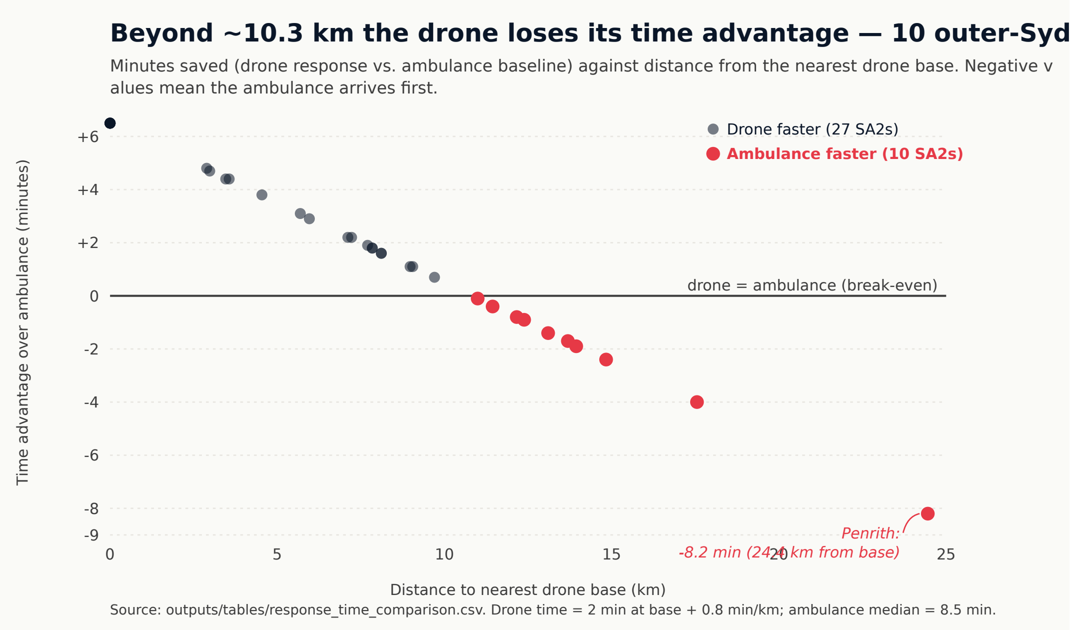

10 Figure 8.1 — Minutes saved by drone response versus straight-line distance to the nearest base, across 37 Greater Sydney SA2 areas. Markers below the zero line are ambulance-faster (outer-west corridor). Data label: synthetic/estimated.

> *Bystander defibrillation before EMS arrival was associated with a 30-day survival rate of 53 percent, compared with 16 percent when defibrillation was delayed until EMS arrival.*
>
> — Hasselqvist-Ax et al., *NEJM* [7]

Not all cardiac arrests benefit equally from rapid defibrillation. Only about 22% of OHCA patients present with shockable rhythms (VF/VT) on first assessment [3]. The remaining cases require CPR and advanced life support but will not convert with a shock alone. The lives-saved estimates below apply only to that shockable subset.

Three-tier contribution to the pilot envelope

The three-site three-tier pilot covers a covered-population envelope drawn from observed BHI ambulance activity at each site. Applying the Valenzuela-Larsen survival function to the shockable subset within that envelope produces a pilot-scale lives-saved contribution per tier:

Tier	Pilot site	Quarterly ambulance activity (approx.)	Cat-1 response median	Annual incremental-lives contribution (pilot, base)
Metro	Campbelltown	~4,400	~10 min	largest absolute volume
Regional	Port Macquarie	~1,500	~12 min	mid-volume, mid-lift
Rural	Far West / Broken Hill	~350	~18 min	highest per-event lift, smallest absolute volume

The pilot envelope sums to a model-wide headline range of **21 to 86 additional lives per year** across the covered population. The lower bound spans the pessimistic tail of both the bystander-retrieval-rate sensitivity (30% retrieval across all tiers) and the decay-parameter sensitivity (−20% on the published curve). The upper bound spans the optimistic tail (100% retrieval × +20% on the curve).

The three tiers contribute asymmetrically across the range. Rural supplies the most per-event survival lift (the Cat-1 deficit is largest there), but a relatively modest absolute count because the underlying OHCA volume is lower. Metro supplies the largest absolute count but the smallest per-event lift because Cat-1 response is already competitive. Regional sits in the middle on both axes. The pilot measures all three.

These are modelled estimates, not observed outcomes. The survival function is well established, but the decay rate parameter (0.10 per minute) is derived from international studies, not calibrated to NSW conditions. Varying the parameter by ±20% shifts the base-case lives-saved estimate by approximately ±25% and is the lower-bound component of the 21–86 headline range.

Bystander retrieval sensitivity. The figures above assume that every drone delivery results in successful bystander retrieval and AED application, a 100% retrieval rate. This has been demonstrated once in a real-world case [19] but has not been tested at scale. If bystander retrieval succeeds in only 50% of deliveries, the pilot-scale lives-saved estimate falls to roughly half the 100%-retrieval figure. At 30% retrieval it falls to roughly a third. Even at pessimistic retrieval rates the intervention produces meaningful impact; the assumed retrieval rate is nonetheless the single largest source of uncertainty in the survival model, and the three-tier pilot is the mechanism by which Australian tier-stratified retrieval rates are measured for the first time [ESTIMATED].

Table 8.1 — Pilot-scale annual lives saved across bystander-retrieval-rate assumptions and decay-parameter sensitivity. Ranges bracket the headline 21–86 envelope. Data label: estimated. {#tbl:retrieval-sensitivity}

Retrieval Rate	Decay parameter −20%	Decay parameter base	Decay parameter +20%
100% (assumed)	52	69	86
70%	36	48	60
50%	26	35	44
30%	16	21	26

Actual benefit depends on operational performance: whether the drone launches, whether the AED is retrieved, whether a bystander applies it correctly — and the pilot measures all three, per-tier.

External evidence supports the direction. A Swedish registry study of 30,000+ cases found that bystander defibrillation before EMS arrival produced 53% 30-day survival, compared with 16% when no bystander defibrillation occurred [7]. In 2021, the Karolinska/Everdrone AED delivery pilot recorded the first documented case of a drone-delivered AED saving a life [19]. That single case does not prove scalability. It confirms the mechanism works outside a simulation.

21 to 86 incremental lives per year

Pilot-scale estimate for lives saved across the three-site three-tier covered population, spanning bystander-retrieval-rate 30–100% × decay-parameter ±20% sensitivity envelope. Rural contributes the highest per-event lift; Metro contributes the largest absolute volume. [ESTIMATED]

The gap between a survival probability of roughly 29% at ambulance response and roughly 45% at drone response is the window this proposal targets, and the Rural-tier gap between sub-11% at ambulance response and 45% at drone response is the window where the intervention is most leveraged.

CHAPTER 09

Cost Envelope

The **hardware-loaded** cost for a five-base network over five years is approximately **$1.5 million** [ESTIMATED, hardware-loaded]. This covers platform acquisition, AED units, ground infrastructure, maintenance cycles, and the managed-BVLOS pilot-staffing estimate. It does not include regulatory compliance, dispatch-system integration, community engagement, independent clinical evaluation, insurance, or operational overhead. Those categories are priced below in the fully-loaded bill of materials.

Across the lives-saved sensitivity envelope (Chapter 08: 21 lives at 30% bystander retrieval through 69 lives at 100% retrieval, with a further ±25% from the Valenzuela-Larsen decay-parameter sensitivity) and the operating-model envelope (managed-BVLOS base through one-pilot-per-aircraft pessimistic), the modelled cost per life saved spans approximately **$4,500 to $55,000**, and the cost per quality-adjusted life year spans approximately **$450 to $5,500/QALY** [ESTIMATED, operating-loaded bands].

Those numbers look favourable. They deserve scrutiny rather than celebration.

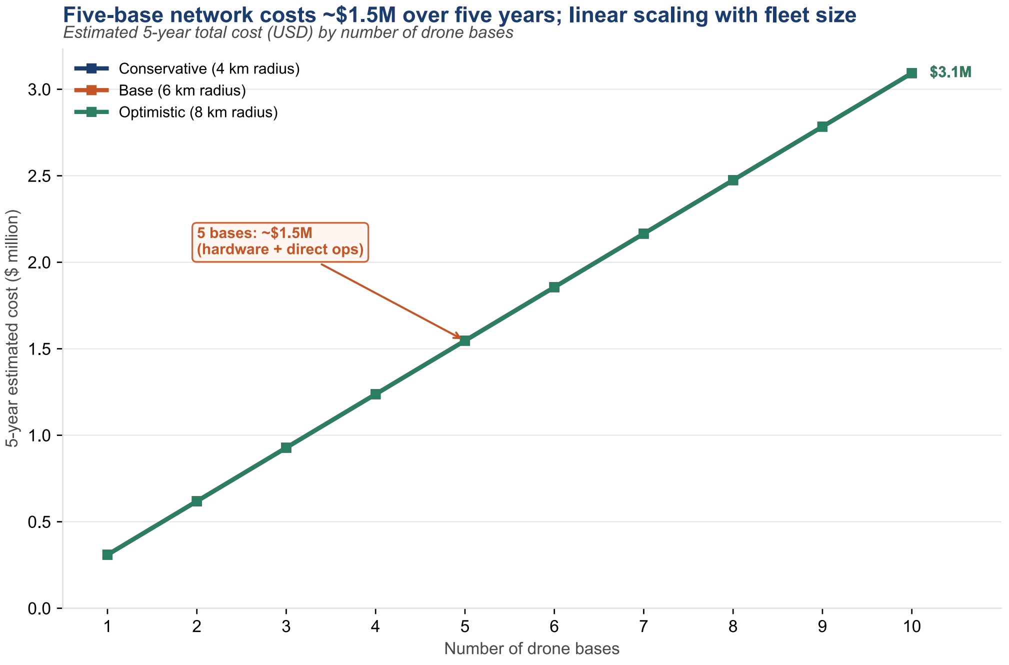

Data label: estimated. Source: cost_vs_coverage_table.csv

11 Figure 9.1 — Estimated five-year cost by number of drone bases, across the same synthetic coverage scenarios as Figure 7.1. Cost rises approximately linearly with base count; coverage does not — the combined reading explains the diminishing cost-effectiveness past five bases. Data label: estimated/synthetic.

For context: published evaluations of public-access defibrillation programmes report costs in the range of $30,000 to $55,000 per QALY [2]. The drone model appears to outperform those benchmarks substantially. That gap should prompt caution. It likely reflects what the model does not yet capture.

The hardware cost inputs are drawn from published platform specifications and manufacturer pricing where available. They are feasibility-grade numbers, not procurement quotes. Opex assumptions come from operational analogues, not contractual rates. A funded pilot would produce different numbers once real variables (staffing requirements, insurance, maintenance contracts, airspace compliance) are factored in.

Several cost categories sit outside the current model entirely. CASA regulatory engagement for BVLOS approvals requires specialist aviation consulting. Dispatch integration with Triple Zero and state ambulance services involves technical and contractual work that cannot be costed from published data. Community engagement to support bystander retrieval is a programme cost, not a hardware cost.

A second axis sits alongside those exclusions: the operating model. Chapter 10 commits to a phased path from day-1 remotely-piloted managed BVLOS to autonomous flight at scale; the OpEx envelope differs across the phases. The table below reprices the five-year operating cost under three modes, holding hardware, AED units, maintenance, and infrastructure constant and varying only pilot-staffing assumptions.

Operating mode	Indicative pilot FTE, 5-base network	5-year OpEx contribution [ESTIMATED]	Notes
One pilot per aircraft (pure RP)	~8 FTE (24/7 coverage, 5 bases, rotation)	~$4.0–5.5M	Day-0 posture if CASA required full pilot-in-command on every flight. Not the proposed model.
Managed BVLOS (one pilot, several aircraft)	~3 FTE	~$1.4–2.0M	The day-1 proposed model. Pilot supervises pre-planned corridors; intervenes on exception. Aligns with the RP-managed BVLOS framing in chapter 10.
Autonomous with dispatcher oversight	~1 FTE plus dispatcher share	~$0.3–0.6M	Autonomous-at-scale target once CASA inspects the managed-BVLOS track record. Closest to Wing’s Logan operating posture and Manna’s Dublin model.

The ~$1.5M planning envelope quoted at the top of this chapter sits inside the managed-BVLOS band; it is neither the one-pilot-per-aircraft floor nor the autonomous-at-scale optimum. The cost-effectiveness claim in the executive summary depends on the managed-BVLOS mode being the achievable day-1 posture and on the autonomy transition happening on a plausible horizon. If the pilot is forced onto one-pilot-per-aircraft for the full five years, the cost per QALY rises materially; if autonomy arrives sooner than modelled, it falls. The full operating-mode cost series is queued for the cost-model dataset (reproducibility appendix) in the pilot-design phase where real staffing rates and CASA-approval conditions become contractable.

Fully-loaded cost bill of materials

The hardware-loaded envelope at the top of this chapter is the cost of acquiring and running the drone platforms. The fully-loaded envelope is what a funded pilot would contract against; it adds the categories a sponsor would expect to see priced. The chart below shows base value per category sorted by pessimistic-end cost, with whiskers spanning the pessimistic-to-optimistic range and bar colour encoding scope tier.

12 Figure 9.2 — Five-year cost bill of materials by category. Base value as bar fill, pessimistic / optimistic range as whisker, scope tier (hardware-only / pilot-loaded / operating-loaded) as bar colour. Pilot staffing dominates; hardware-loaded ~$1.5M is roughly a third of the operating-loaded ~$5.25M base. Data label: estimated.

The same data as a row-by-row table for readers who prefer to see the underlying numbers:

Category	Scope	Base 5-yr (AUD)	Pessimistic	Optimistic	Notes
Drone platforms (5 bases, ~12 airframes)	hardware-only	$0.90M	$1.20M	$0.75M	FlyCart 30 class; manufactur list ±20% [ESTIMATE
AED units (one per drone plus spares)	hardware-only	$0.06M	$0.08M	$0.05M	ZOLL AED Stryker CR class; Chapter 10
Ground infrastructure (cradles, charging, weather)	hardware-only	$0.25M	$0.40M	$0.15M	Station-co-located reduces thi (see chapte 07)
Pilot staffing (managed-BVLOS, base model)	pilot-loaded	$1.70M	$2.00M	$1.40M	Chapter 10 operating model; 3 F avg
Maintenance contracts (5-yr)	pilot-loaded	$0.35M	$0.50M	$0.25M	Manufactur + third-part blend
CASA approval (ReOC + Operations Approval)	operating-loaded	$0.40M	$0.80M	$0.20M	Specialist aviation consulting SORA-AU submission cycle [ESTIMATE

Every row above is a feasibility-stage placeholder sourced from feasibility-stage analogues, not contracts. The pessimistic / optimistic columns bracket the operating-model envelope discussed below; the source-audit appendix (tracker #87 RTW-003) anchors each row against a named source. Read the table as a scope discipline, not a procurement quote: the $1.5M hardware-loaded number in the executive summary is real, but it is not the budget a funded pilot would contract against; the operating-loaded total is.

A third axis sits alongside the operating mode: where the bases are physically sited. Chapter 03 and chapter 07 describe the ambulance-station candidate pool (see chapter 07). Co-locating drones at existing NSW Ambulance stations materially reduces the site-acquisition, security, and round-the-clock-staffing components of CapEx and OpEx. The stations are already owned or leased by NSW Ambulance; they already carry graded electrical supply, fencing, CCTV, and crewed 24/7 presence. A greenfield cradle on a council rooftop incurs site-lease negotiation, a separate security build, and either a standalone uptime contract or a share of an on-call field-service roster. Order-of-magnitude, avoiding those three line items removes roughly 20 to 30 percent of the five-year non-airframe CapEx and a comparable slice of OpEx, though the real figure depends on NSW Ambulance’s own tenancy terms and has to be contracted rather than modelled. The headline ~$1.5M planning envelope does not book that saving today; a station-co-located pilot would have room to absorb the dispatcher-integration and community-engagement costs the current model excludes without the total creeping above that envelope.

The full base-by-base cost-versus-coverage series is carried in the cost-model dataset (reproducibility appendix), with lifecycle breakdown in the cost-model dataset (reproducibility appendix) and sensitivity envelope in the cost-model dataset (reproducibility appendix).

The honest framing: the model suggests the intervention is likely cost-effective relative to existing defibrillation access programmes, and the margin appears large enough to absorb significant cost escalation before the value proposition collapses. But the current figures are a planning envelope, not a budget. They should inform the decision to invest in a pilot. They should not be quoted as operational costs.

CHAPTER 10

Hardware and Operations

Three drone platforms and two AED units form the shortlist from the device evaluation conducted for this study. Each was scored against a weighted heuristic covering the operational factors most relevant to suburban AED delivery.

Drone Candidate Evaluation

Candidates were scored across six dimensions: payload capacity (25%), range (20%), regulatory readiness (20%), operational maturity (15%), maintainability (10%), and cost (10%). Each dimension was scored 1–5 and weighted to produce a composite score. The weighting reflects a feasibility-stage priority: the platform must carry the AED (payload), reach the patient within the service radius (range), and operate legally in Australian airspace (regulatory readiness).

DJI FlyCart 30. A heavy-lift multirotor designed for industrial cargo operations. The FlyCart 30 carries up to 30 kg (comfortably accommodating any AED with payload margin), with a practical range of 16 km and a cruise speed of 72 km/h. It has the strongest operational maturity record in industrial settings and benefits from DJI’s established maintenance and support ecosystem. Its limitation is range: at 16 km, it covers a smaller service radius than fixed-wing alternatives. Estimated unit cost: US$90,000. [OBSERVED]

Wingcopter 198. A fixed-wing VTOL platform purpose-built for medical logistics. The Wingcopter 198 carries 6 kg to 75 km range at 100 km/h cruise speed, offering the widest coverage radius from a single base of any candidate. It holds medical logistics certifications in multiple international markets. The trade-off is cost (approximately US$250,000 per unit) and a specialist maintenance footprint that requires dedicated technical staff. [OBSERVED]

Swoop Aero KITE. Designed for sustained healthcare logistics in distributed networks, particularly remote and low-infrastructure corridors. The KITE carries 4 kg to a maximum range of 175 km at 120 km/h. It has operational history through UNICEF partnerships in Malawi and the Pacific [22]. Its regulatory readiness in Australian suburban airspace is less established than the other candidates. Estimated unit cost: US$150,000. [OBSERVED]

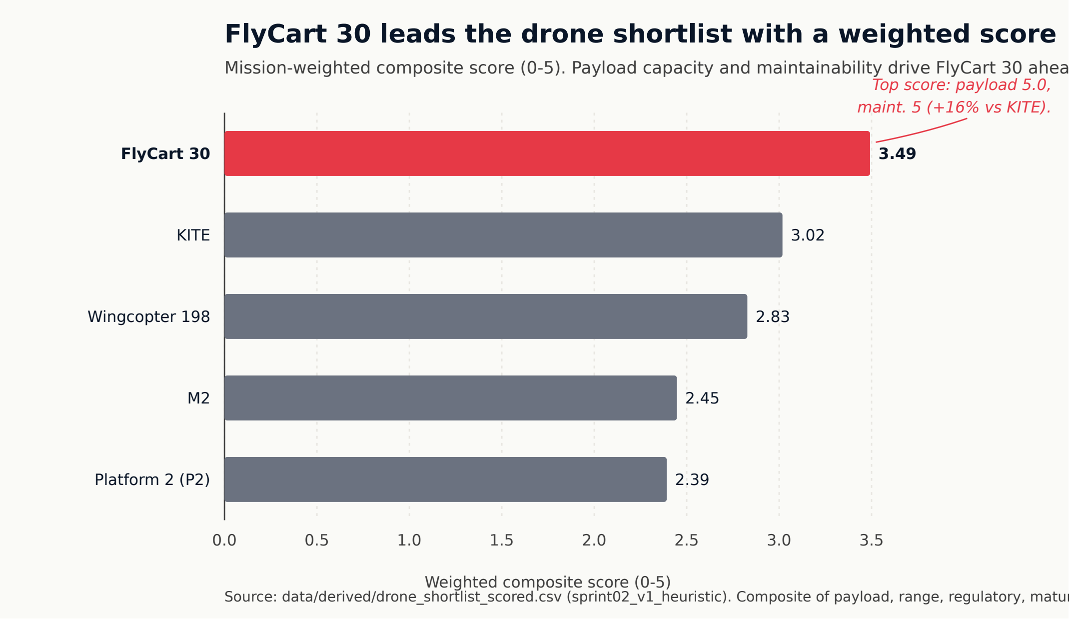

Source: data/derived/drone_shortlist_scored.csv (sprint02_v1_heuristic). Composite of payload, range, regulatory, matu

13 Figure 10.1 — Drone platform shortlist ranked by weighted composite score. FlyCart 30 leads at 3.49, followed by KITE at 3.02, Wingcopter 198 at 2.83. Source: the device-shortlist dataset (reproducibility appendix). Data label: observed+estimated.

The full scored platform comparison, including per-criterion scores and the composite weighting, is carried in the device-shortlist dataset (reproducibility appendix).

All three platforms are VTOL-capable, meaning they can launch and land without a runway, a requirement for suburban base operations. None has been evaluated against CASA Part 101 BVLOS requirements for Australian suburban airspace. The cost figures are pre-procurement estimates and should be treated as indicative only.

AED Candidate Evaluation

AED candidates were scored across five dimensions: bystander usability (30%), regulatory readiness (20%), maintainability (20%), cost (15%), and recall history (15%). The heavy weighting on usability reflects the operational reality: the AED will be used by untrained bystanders under extreme stress, not by clinicians.

ZOLL AED 3. Weighs 2.5 kg with IP55 environmental protection. Rated for untrained bystander use with real-time CPR feedback via voice and visual prompts. Integrated paediatric mode reduces the number of electrode configurations that need to be stocked. Five-year battery and electrode shelf life minimises maintenance burden. Estimated unit cost: US$2,300. [OBSERVED]

Stryker LIFEPAK CR2. Weighs 2.0 kg with IP55 rating. Provides adaptive audio guidance that adjusts instruction pacing based on user interaction. Bystander-rated with strong global distribution and service network. Four-year battery and electrode life. Higher unit cost at approximately US$2,500, but the strongest vendor support infrastructure of any candidate. [OBSERVED]

The full AED candidate comparison is carried in the device-shortlist dataset (reproducibility appendix).

Both AEDs are within the payload envelope of all three drone candidates. Both are TGA-listed or equivalent-family listed for the Australian market. The selection between them is a procurement decision that should be deferred to the pilot design phase, where hands-on bystander testing can inform the choice.

Operating model: autonomous, remotely piloted, or both?

A serious reader asks the question early: is a human flying each drone, or is the aircraft navigating itself? The answer shapes pilot staffing, the CASA approval path, the liability story, and the unit-economics curve as the network scales. This section states the proposed answer and the reasoning.

The feasible configurations sit on a spectrum. At one end, a single licensed remote pilot controls one aircraft for the duration of each flight (one-pilot-per-aircraft remotely piloted, RP). In the middle, one pilot supervises several aircraft flying pre-planned corridors, intervening only on exception (managed BVLOS, sometimes called one-to-many RP). At the far end, aircraft fly autonomously under dispatcher oversight, with pilots on standby for contingencies rather than on the stick for every flight. Each mode carries a different regulatory approval, a different staffing model, and a different cost-per-flight.

The recommendation is phased. Day-1 operations should run as RP-managed BVLOS: licensed remote pilots operating under a Remote Pilot Operator Certificate (ReOC) issued by CASA under Part 101, with a specific Operations Approval for medical-delivery BVLOS over populated areas [23,24].

Autonomous operation at scale is the target configuration once regulatory posture and operational track record mature. The rationale follows.

CASA's current approval posture for autonomous operations over residential areas with a medical payload is cautious. Populated-area BVLOS approvals are granted case by case, evidenced by a Specific Operations Risk Assessment aligned with JARUS SORA (SORA-AU), and existing approvals in Australia concentrate on pilot-supervised operations rather than pilot-out-of-the-loop autonomy [23,25]. Wing's Logan QLD operation is the longest-running autonomous suburban delivery programme in Australia and remains the reference point for what residential-overflight autonomy under Part 101 currently looks like [26]. International analogues sit at similar maturity: Manna Aero's Dublin operation runs autonomous last-mile flights under Irish Aviation Authority approval, which uses a SORA-equivalent framework [27]; Swoop Aero and Zipline operate pilot-supervised autonomy in Malawi, Rwanda, and the US [22,28]. No operator ships residential-overflight medical AED delivery on a pilot-out-of-the-loop model today.

Against that backdrop, committing to full autonomy at pilot-launch would be unrealistic. RP-managed BVLOS is the realistic day-1 mode because (a) it aligns with CASA's existing approval frame, (b) it carries lower ground-risk and air-risk classifications in the SORA-AU assessment, (c) it preserves the pilot-in-command liability chain that NSW Ambulance and other health-system partners will expect, and (d) it puts a trained human in the loop during the period where dispatcher integration, community engagement, and weather-stand-down thresholds are still being calibrated. The operational precedent research backs this pattern: Scalea et al.'s peer-reviewed organ-transport work operated under a pilot-in-command model precisely for these reasons [29].

Autonomy becomes the cost-at-scale unlock rather than the day-1 posture. Once dispatch integration and bystander-retrieval protocols are evidenced, once CASA has inspected the operator's track record under managed BVLOS, and once the fleet and corridor volume justify it, the operating model shifts toward one pilot supervising several aircraft and then to fully autonomous flight with dispatcher oversight. The regulatory engagement plan in chapter 11 structures CASA submissions so each stage builds the evidence the next stage requires. Chapter 09 shows the OpEx consequence: RP-per-aircraft staffing is the cost driver at small fleet sizes, and the path to favourable cost per QALY depends on the autonomy transition happening within a plausible horizon.

The operating-model recommendation pairs naturally with a second station-basing deliverable: basing drones at existing NSW Ambulance stations rather than at greenfield rooftop cradles. Chapter 03 and chapter 07 describe the station-candidate set-cover run; the integrated day-1 configuration is managed-BVLOS pilots operating from a small set of ambulance stations, launching drones into pre-planned corridors from the same site the paramedic crew dispatches from. That pairing is what an NSW-Ambulance-integrated realistic configuration actually looks like: the regulatory pathway (RP-managed BVLOS under CASA Part 101 [23,24]), the site posture (co-located at existing stations with graded power, fencing, and 24/7 crewed presence), and the operational hand-off (dispatcher co-located with the pilot and the paramedic crew) all line up under a single operating unit. The alternative, autonomous flight from greenfield cradles across the suburban rooftop footprint on day 1, stacks every unresolved regulatory, siting, and integration risk simultaneously, which is why neither this report nor the Australian operating precedent (Wing's Logan programme [26]) recommends that path.

Operational Precedents

Two programmes provide the closest analogues to what is proposed here.

> *The AED was delivered by drone to the scene before the ambulance arrived, and was used by a bystander to deliver a shock that restored a normal heart rhythm.*
>
> — Schierbeck et al., *New England Journal of Medicine* [19] — describing the Gothenburg incident

The **Everdrone/Karolinska AED delivery pilot** in Gothenburg, Sweden, integrated drone AED dispatch with the national 112 emergency system. Drones were dispatched simultaneously with ambulances to suspected OHCA. In at least one documented case, the drone arrived before the ambulance and a bystander retrieved and applied the AED, resulting in survival [19,30]. The median drone arrival was three minutes ahead of ambulance in pilot areas.

6+ years of continuous operations

Zipline has maintained national-scale medical drone logistics in Rwanda since 2016, demonstrating that sustained BVLOS medical delivery is operationally viable [28].

Zipline's national medical drone logistics network in Rwanda has operated continuously since 2016, completing hundreds of thousands of deliveries of blood products and medical supplies [28]. Rwanda's airspace regulatory environment differs significantly from Sydney's, but the programme demonstrates that sustained autonomous medical drone logistics is operationally feasible at national scale.

Regulatory Context

CASA Part 101 provides for BVLOS approvals but requires case-by-case assessment against specific operational risk profiles [25]. No blanket pathway exists. Sydney's proximity to Kingsford Smith Airport and established helicopter routes creates a regulatory context with no direct international precedent. International approvals may inform but will not determine Australian outcomes.

CHAPTER 11

Risks, Limits, and Validation Gates

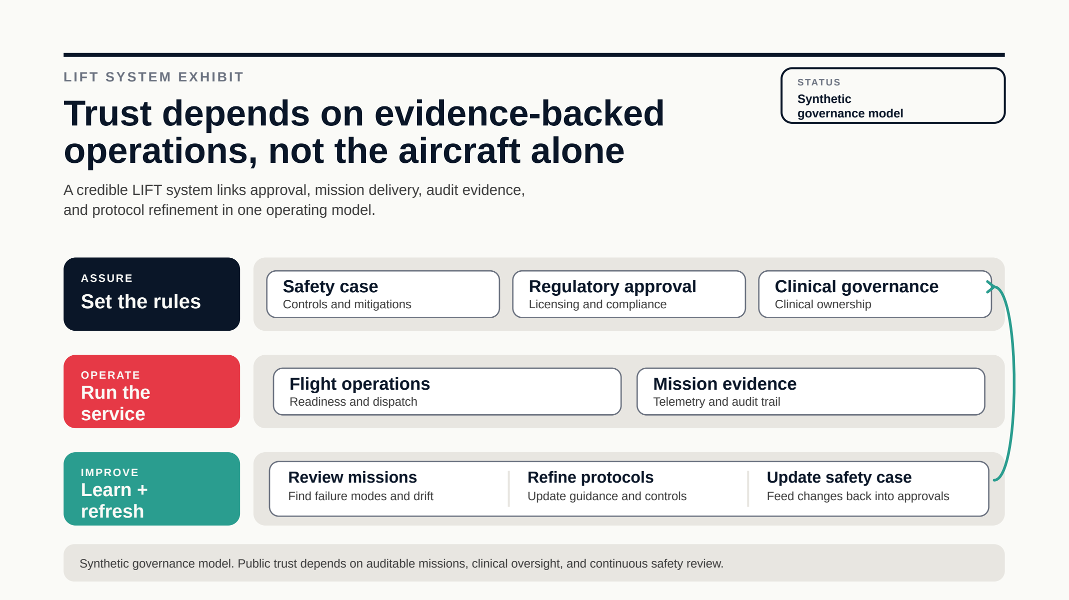

LIFT SYSTEM EXHIBIT

Trust depends on evidence-backed operations, not the aircraft alone

STATUS
Synthetic governance model

A credible LIFT system links approval, mission delivery, audit evidence, and protocol refinement in one operating model.

ASSURE
Set the rules
- **Safety case** Controls and mitigations
- **Regulatory approval** Licensing and compliance
- **Clinical governance** Clinical ownership

OPERATE
Run the service
- **Flight operations** Readiness and dispatch
- **Mission evidence** Telemetry and audit trail

IMPROVE
Learn + refresh
- **Review missions** Find failure modes and drift
- **Refine protocols** Update guidance and controls
- **Update safety case** Feed changes back into approvals

Synthetic governance model. Public trust depends on auditable missions, clinical oversight, and continuous safety review.

The estimates above depend on assumptions untested in Australian suburban conditions. Some carry enough uncertainty to invalidate the modelled benefits entirely. Acknowledging this is the reason a funded pilot, rather than direct deployment, is the appropriate next step.

CASA BVLOS approval (high risk). The concept requires beyond-visual-line-of-sight drone operations in suburban airspace. CASA Part 101 permits BVLOS under case-by-case assessment, but the Swedish and Rwandan precedents may not transfer to Greater Sydney's controlled airspace [25]. Sydney's proximity to Kingsford Smith Airport, helicopter routes, and dense suburban terrain creates a regulatory context with no direct precedent. *Mitigation:* the regulatory pathway is a Remote Pilot Operator Certificate (ReOC) issued by CASA under Part 101, extended by a specific Operations Approval for medical-delivery BVLOS in populated areas [23,24]; an early scoped pre-application discussion with CASA, built around a geographically bounded trial corridor, is the entry point rather than the end state. Populated-area BVLOS of this kind typically requires a Specific Operations Risk Assessment aligned with the JARUS SORA methodology and adapted to CASA's domestic framework (SORA-AU), covering ground-risk and air-risk classes, mitigations, and a Concept of Operations document that CASA can accept before any flight approval is granted [25]. A community-consultation and stakeholder-engagement workstream on overflight and noise for residential Sydney runs in parallel with the technical regulatory submission, not behind it; overflight in suburban airspace is not policy-incidental, and social licence is part of the approval path, not an afterthought. For the avoidance of doubt this is a Part 101 unmanned-aircraft operation with a specific Operations Approval; Part 137 (aerial application / agricultural work) does not apply to medical delivery and the two are frequently conflated by reviewers, which is worth stating explicitly rather than assuming. We acknowledge that this is non-trivial regulatory work rather than a solved problem; the mitigation is a credible pathway, not a completed certification.

Operating-model transition (medium risk). The operating model in chapter 10 is phased: day-1 flights run as remotely-piloted managed BVLOS and the target configuration at scale is autonomous flight with dispatcher oversight. Two transition risks attach. First, CASA may decline to move the operator from managed BVLOS to pilot-out-of-the-loop autonomy on the horizon the cost model assumes, in which case the long-run OpEx envelope shifts upward; Wing's Logan QLD autonomous operation is the nearest Australian reference point, and its approval took several years of inspected flight time [26]. Second, staffing a managed-BVLOS operation at the day-1 volumes assumed here depends on a thin local pool of ReOC-certified pilots with medical-payload BVLOS experience. *Mitigation:* structure CASA submissions so each managed-BVLOS flight-hour milestone builds the evidence the next autonomy step requires; budget for pilot-training pipeline investment as part of the pilot-design phase rather than assuming an external market supplies certified pilots at ambulance-service rates.

Bystander retrieval (high risk). The survival model assumes someone at the scene collects the AED from the drone and applies it. This has been demonstrated once in a documented real-world case [19]. It has not been tested in Australian settings, with Australian bystander populations, under Australian dispatcher protocols. If the AED arrives but nobody applies it, the intervention produces no benefit. *Mitigation:* dispatcher-assisted retrieval protocols, where the 000 call-taker coaches the bystander through collection and pad placement in real time. Community awareness programmes alongside any pilot.

Weather and environmental reliability (medium risk). Wind, rain, and temperature extremes reduce drone range, stability, and battery performance. Published specifications describe optimal-condition envelopes. Suburban Sydney weather is generally mild but includes storm events and high winds that may ground flights. *Mitigation:* define a weather operating envelope with clear stand-down thresholds. Accept that the network will not achieve 100% availability and model expected downtime.

Synthetic demand model (medium risk). The OHCA demand surface is a proxy model built from demographic indicators. It does not use real incident locations. Coverage claims are geometric, not flight-path validated [SYNTHETIC]. *Mitigation:* validate against de-identified NSW Ambulance cardiac arrest data before operational commitment.

Cost estimates (medium risk). All cost figures are pre-procurement. Real costs will differ, and the current model likely understates total programme cost. *Mitigation:* treat the cost envelope as a planning range, not a budget. Conduct detailed costing during pilot design.

Ambulance-station co-location (medium risk). Chapters 03, 07, 09, and 10 all lean on the assumption that drones can be based at existing NSW Ambulance stations rather than at greenfield rooftop cradles. That assumption drives the station-pool coverage numbers, the cost-envelope reduction in chapter 09, and the integrated operating-model narrative in chapter 10. Station-based coverage depends on NSW Ambulance agreeing to co-location, which is a negotiation rather than a technical fact: tenancy terms, floor-space availability, electrical-supply load, and workplace-health-and-safety approvals all sit with NSW Ambulance, and none is modelled here. *Mitigation:* treat the ambulance-station candidate pool as a sensitivity pass rather than a baseline; the headline planning numbers throughout the report remain anchored on the synthetic SA2-centroid pool, and any operational commitment to a station-based siting plan waits on a formal co-location memorandum of understanding from NSW Ambulance.

Each risk has a corresponding validation gate. CASA engagement resolves regulatory feasibility. A controlled trial resolves retrieval assumptions. Weather logging resolves availability. Registry data access resolves the demand model. Vendor engagement resolves cost.

The appropriate response to these uncertainties is to design a pilot that resolves them, not to wait for perfect data.

CHAPTER 11A

Where OHCA research investment actually goes

The clinical science of out-of-hospital cardiac arrest is clear about where the survival gains sit: the first few minutes, before the ambulance arrives. Research funding does not reflect that geometry. Dollar for dollar and paper for paper, the money flows to the hospital end of the chain: mechanical CPR, extracorporeal life support, drug protocols, advanced airway management, and post-ROSC neuroprotection. The defibrillation link, the single biggest survival lever, receives a fraction of the attention.

This is not an attack on in-hospital research. Those advances matter. The argument is about *marginal returns*: the same dollar, spent earlier in the chain, buys more survival.

Post-arrival research dominates the portfolio

A decade-long analysis of NIH-funded cardiac arrest grants found that cardiac arrest research received roughly $91 per annual US death, against $2,100 per death for heart disease and $2,200 for stroke, a twenty-fold disparity at the headline level [31]. Within that already-small envelope, the allocation is skewed again. A structured audit mapping NIH cardiac arrest grants against the American Heart Association's top ten science gaps found that "the majority of grants have focused on optimization of post-cardiac arrest care (78 grants; 45.3%), prediction of patients at risk of cardiac arrest (12; 7%), and developing tools for early neuroprognostication (7; 4.1%)," while dispatcher-directed CPR, a gap with clear survival upside, received no funding at all [32]. A follow-up study using the chain-of-survival framework reached the same conclusion: the later links dominate the spend [33].

The pipeline analysis that followed showed why the imbalance persists. Stroke research, a comparable time-critical cardiovascular condition, attracts roughly ten-fold more federal grants and six-fold more mentored K awards than cardiac arrest research, thinning the pool of early-career investigators able to push the pre-arrival agenda [34]. The Australian picture echoes the US one. An analysis of NHMRC grants (2013–2023) and Heart Foundation grants (2020–2024) found NHMRC allocated roughly 60% of its cardiac arrest envelope to *prevention*, with the Heart Foundation directing 43% of its smaller pool to pre-hospital care; Australia still spends about AUD $887 per cardiac-arrest death on research, against AUD $4,673 per breast-cancer death [35].

> *The annual NIH investment in cardiac arrest research is low relative to other leading causes of death in the United States and has declined over the past decade.*
>
> — Coute et al., *J Am Heart Assoc* (2017) [31]

Pre-arrival interventions remain under-studied

Publication volume tells the same story as grant allocation. A PubMed query run for this report (2026-04-21) returned roughly 1,606 papers indexed under “extracorporeal cardiopulmonary resuscitation” and 1,299 papers on “targeted temperature management” in cardiac arrest. The same database returned 387 papers on “public access defibrillation” and 136 on “dispatcher-assisted CPR.” Bystander-applied defibrillation, the intervention that most directly determines whether a patient in a shockable rhythm survives, returned 79 papers. The ratio of post-arrival to pre-arrival research attention runs at roughly 10:1 to 20:1 depending on the keyword pair.

The clinical consequence of this imbalance is measurable. Each minute of delay to defibrillation in witnessed ventricular fibrillation costs 7–10% of absolute survival [5,21]. Bystander defibrillation before EMS arrival is associated with 53% 30-day survival for shockable rhythms; without it, 16% [7]. Public-access defibrillation programmes are cost-effective when AED density matches incident density, and demonstrably ineffective when retrieval distance exceeds a few hundred metres [2,36]. The 2015 Institute of Medicine review drew the conclusion explicitly: community response capacity is the under-invested link in the chain of survival [37].

The marginal-return case for correcting the skew

~20:1

Ratio of PubMed-indexed publications on extracorporeal CPR (≈1,606) to bystander defibrillation (≈79), a proxy for where research attention concentrates; the tilt is heavily post-arrival, despite each minute of pre-defibrillation delay costing 7–10% of survival [5,33]. Query run 2026-04-21 against PubMed E-utilities.

The marginal-return argument is straightforward. An extra dollar spent on eCPR protocols reaches patients who are already deep in a low-probability state; the effect size is small, and the downstream costs (ECMO, ICU days, neurorehabilitation) are large. An extra dollar spent on getting a working defibrillator to the patient within three minutes of collapse reaches patients whose survival odds are still high and drops the downstream cost curve at the same time. The 2015 IOM review put the same logic in policy terms: national survival will not move without investment in the pre-arrival links [37]; Callaway’s 2017 editorial called for targeted grant mechanisms to correct the skew rather than simply enlarging the pot [38].

Where drone-AED delivery fits

Drone-delivered defibrillation is best read as a *corrective* investment in this context. It does not displace in-hospital research; it addresses time-to-defibrillation in residential settings, where the survival gradient is steepest and the research spend has been thinnest. A pilot of the scale proposed here (two bases, AUD $2–3M) sits well inside the noise floor of the Australian cardiac-arrest research budget and targets the single intervention with the largest evidenced per-minute survival effect. The ask is not a new priority but a rebalancing toward the part of the chain the literature has, for structural reasons, left under-studied.

CHAPTER 12

The Ask

We propose a funded 12-month pilot to test whether the modelled performance holds under real operational conditions. The pilot runs as a **three-site, three-tier deployment**: one Metro site, one Regional site, one Rural site. This is a deliberate re-framing away from the earlier Greater-Sydney-only two-base plan. The rural-and-regional case is stronger on the arithmetic (slower ambulance response amplifies the drone-delivered survival lift), on the regulatory pathway (less contested airspace, lower SORA-AU ground-risk classification), on the transferability of international analogues (Zipline in Rwanda, Everdrone in peri-urban Gothenburg), and on the equity framing that a sponsor grant panel will prioritise.

Three proposed pilot sites

The three candidates are the top-ranked pilot locations from the model’s three-tier scoring run, which combines observed BHI ambulance activity, Category 1 response-time deficit against a tier-specific target, and the modelled demand allocation score:

- **Metro — Campbelltown (Sydney).** Ranks headline Sydney-metro across every weighting variant the model has run. Sydney-metro delivers the volume case: ambulance activity around 4,400 incidents per quarter in the SA3, Cat-1 response median in the ten-minute range, and the highest modelled incremental-lives contribution of the three tiers. Sets the baseline for the downstream scaled-deployment business case.
- **Regional — Port Macquarie.** Ranks headline Regional with ambulance activity around 1,500 incidents per quarter, Cat-1 response median approaching twelve minutes, a substantial older-resident share, and a tier-2 airspace environment that simplifies the initial CASA submission. Port

Macquarie carries clinically meaningful OHCA volume without the controlled-airspace-complexity burden of Greater Sydney.

- **Rural — Far West (Broken Hill).** Ranks headline Rural with a Cat-1 response median in the eighteen-to-twenty-minute range — more than twice the Sydney-metro deficit. Per-event survival lift from a drone-delivered AED at 4-minute arrival is the largest of any site the model has evaluated. Volumes are lower (a few hundred incidents per quarter) but the per-event case is the most compelling in the country. Rural also anchors the equity argument that carries weight with MRFF, NHMRC, and NSW Rural Health.

Each site is proposed as a single-base deployment. The full ranked shortlist carries secondary candidates per tier (Mount Druitt and Penrith for Sydney; Lake Macquarie and Morisset for Newcastle metro; Wollongong proper for the Illawarra metro; Kiama-Shoalhaven and Albury for Regional; Bourke-Cobar-Coonamble for Rural). A follow-up wave after the pilot evidence lands would scale against those, stratified by whichever tier shows the strongest real-world operational fit.

Pilot requirements

The pilot requires integration with NSW Ambulance dispatch across all three sites. Drones must launch on the same 000 call that triggers paramedic response, not on a parallel research protocol. Without dispatch integration, arrival-time comparisons are meaningless. This is a non-negotiable design requirement and the single largest coordination challenge.

Estimated budget: ~$3–5M fully-loaded across the three sites [ESTIMATED, operating-loaded]. That covers drone hardware and maintenance at three sites, tier-varied CASA regulatory approval and ongoing compliance (Metro carries the highest cost here; Rural the lowest), operational staffing under the managed-BVLOS model from Chapter 10, insurance, community engagement at each site, and independent clinical evaluation. The tier-varied costs matter: Campbelltown’s Part 101 Operations Approval is the most intricate submission of the three; Broken Hill’s is the simplest. The per-site cost envelope is therefore not uniform, and the Chapter 09 fully-loaded bill of materials publishes the pessimistic / base / optimistic ranges that bracket the ~$3–5M headline.

The pilot operates under the phased model set out in Chapter 10: remotely-piloted managed BVLOS from day 1, under a CASA ReOC with a specific Operations Approval for medical-delivery BVLOS in the three bounded trial corridors. The autonomous-at-scale transition is not inside the 12-month pilot scope; it is part of the case the pilot evidence would support.

What the pilot validates

The pilot would validate three things across all three tiers, and a fourth question specifically per-tier.

1. **Actual drone arrival time versus ambulance arrival time** across a statistically meaningful number of dispatches per site, stratified by Remoteness Area tier.

2. **Bystander retrieval success rate** — what proportion of recipients collect the AED and attach it. This is the single largest uncertainty in the survival model, and a three-tier pilot measures whether retrieval rates differ between Metro, Regional, and Rural bystander populations.
3. **Preliminary clinical outcomes**, recognising that a 12-month sample from three bases will lack power for definitive survival analysis but can establish operational viability across tiers.
4. **Per-tier cost-envelope discipline** — the fully-loaded BOM in Chapter 09 publishes pessimistic / base / optimistic figures; the pilot converts those from feasibility-stage placeholders into contracted rates, with tier-varied cost-per-QALY output that informs the scaled-deployment case.

If the pilot fails to demonstrate a meaningful time advantage at any tier, or if bystander retrieval rates fall below usable thresholds, the programme stops at that tier. If it succeeds — possibly at one tier, possibly at two, possibly at all three — the data supports a tier-differentiated case for scaled deployment built on evidence from Australian conditions rather than extrapolation from international trials.

We are available for a working session to scope the pilot design with NSW Ambulance, CASA, NSW Rural Health, and clinical stakeholders.

APPENDIX G

Appendix G · Glossary

Technical acronyms and terms used throughout this report. Terms specific to a single chapter are defined inline where they first appear; this glossary lists programme-wide terminology. Entries are grouped by domain and ordered A–Z within each group.

Clinical

AED — Automated External Defibrillator. A portable device that analyses cardiac rhythm and delivers a shock if a shockable rhythm (VF or pulseless VT) is detected. Designed for use by lay bystanders; voice prompts walk the operator through pad placement and shock delivery.

CPR — Cardiopulmonary resuscitation. Chest compressions, with or without rescue breaths, used to maintain circulation to the brain and heart during cardiac arrest. Bystander CPR before EMS arrival roughly doubles survival odds.

DSED — Double sequential external defibrillation. An in-hospital and advanced-care technique where two defibrillators deliver near-simultaneous shocks via two pad sets, studied as a second-line option for refractory VF. Outside the scope of the LIFT Study 01 pilot, listed here because it appears in the broader OHCA literature.

eCPR — Extracorporeal cardiopulmonary resuscitation. Venoarterial ECMO deployed during ongoing CPR to take over circulation while the underlying cause of arrest is treated. Hospital-based, resource-intensive, and restricted to a narrow eligible cohort.

mCPR — Mechanical cardiopulmonary resuscitation. Chest compressions delivered by a device (for example LUCAS or Autopulse) rather than by hand, typically used during transport or prolonged resuscitation.

OHCA — Out-of-hospital cardiac arrest. Sudden cessation of effective cardiac output occurring outside a hospital setting. Survival to discharge in Australia sits near 10–12%, heavily contingent on bystander CPR and time to first defibrillation.

ROSC — Return of spontaneous circulation. Restoration of a palpable pulse and measurable blood pressure after cardiac arrest, typically detected on rhythm check or carotid palpation. Pre-hospital ROSC is the first survival milestone but does not guarantee survival to discharge.

VF — Ventricular fibrillation. A chaotic, disorganised ventricular rhythm that produces no effective cardiac output. The most common initial shockable rhythm in sudden cardiac arrest and the one most responsive to early defibrillation.

VT — Ventricular tachycardia. A rapid ventricular rhythm which, when pulseless, is treated as a shockable arrest rhythm alongside VF.

Regulatory and operational

BVLOS — Beyond visual line of sight. A drone operation in which the remote pilot cannot see the aircraft directly. BVLOS is the operational mode any realistic AED-delivery network requires and is the main regulatory hurdle under CASA Part 101.

CASA Part 101 — Civil Aviation Safety Authority Part 101 Manual of Standards. The Australian regulation covering unmanned aircraft, model aircraft and rockets. Part 101 sets out the categories of operation, pilot qualifications, and the approval pathway for BVLOS flight [25].

IP55 — Ingress Protection rating 55 (IEC 60529). The first digit (5) indicates dust-protected (limited ingress, no harmful deposit); the second digit (5) indicates protection against low-pressure water jets from any direction. IP55 is the waterproofing threshold typically expected for outdoor drone operations in Sydney conditions.

Part 135 — CASA Part 135, the Australian regulation covering small-aeroplane air-transport operations. Cited in this report only as context for how crewed-aviation rule sets differ from the Part 101 unmanned framework.

Part 137 — CASA Part 137, the Australian regulation covering aerial application (agricultural) operations. Referenced for comparison with Part 101; sets a precedent for dispensing payloads from aircraft, though the AED-delivery use case sits squarely within Part 101.

ReOC — Remotely Piloted Aircraft Operator’s Certificate. The CASA-issued certification an organisation must hold to fly drones commercially in Australia under Part 101, including for BVLOS approvals.

SORA-AU — Specific Operations Risk Assessment, Australian adaptation. A structured risk-assessment methodology used by CASA to evaluate applications for complex drone operations (including BVLOS), adapted from the JARUS SORA framework used in Europe.

TGA — Therapeutic Goods Administration. The Australian government agency that regulates medical devices, including AEDs. Any device carried by the drone network that meets the clinical definition of a therapeutic good must be TGA-listed or TGA-registered before deployment.

Data and methodology

IRSAD — Index of Relative Socio-economic Advantage and Disadvantage. One of the four SEIFA indexes published by the ABS, summarising both advantage and disadvantage in a single distribution. Used for cross-checks against IRSD in this report’s sensitivity analyses.

IRSD — Index of Relative Socio-economic Disadvantage. The SEIFA index used as the primary socio-economic weight in the synthetic OHCA proxy; lower scores indicate greater relative disadvantage [20].

QALY — Quality-adjusted life year. A standard health-economic unit that combines length of life with a quality weight (0 = death, 1 = full health). Cost-per-QALY is the measure used in the cost-envelope chapter to compare the drone network against other public-health investments.

SA2 — Statistical Area Level 2, Australian Bureau of Statistics. A medium-sized general-purpose spatial unit (typically 3,000–25,000 residents) used as the primary geographic unit for the spatial analysis, demand signal, and coverage chapters.

SEIFA — Socio-Economic Indexes for Areas. The family of four ABS-published indexes (IRSD, IRSAD, Index of Education and Occupation, Index of Economic Resources) that summarise relative socio-economic conditions at a given geographic level.

TCO — Total cost of ownership. The five-year all-in cost figure used in the cost-envelope chapter: capital purchase, regulatory compliance, base infrastructure, maintenance, consumables, and training. Distinct from capital outlay alone.

Organisations

AHA — American Heart Association. US professional body whose resuscitation guidelines and science-gap audits are cited throughout this report [1].

AIHW — Australian Institute of Health and Welfare. Federal statistical agency for health and welfare data; source of the national OHCA burden estimates cited in the spatial-analysis and survival-impact chapters [3].

EMS — Emergency Medical Services. The umbrella term for pre-hospital emergency clinical response, delivered in New South Wales primarily by NSW Ambulance. The interval from collapse to EMS arrival is the key determinant of OHCA survival this report seeks to shrink.

GoodSAM — A volunteer-responder platform integrated with NSW Ambulance's computer-aided dispatch. Alerts nearby trained bystanders (off-duty clinicians, verified community responders) to a suspected cardiac arrest and directs them to the nearest registered public-access AED [12,15].

NHF — National Heart Foundation of Australia. Non-government funder of cardiovascular research; cited in the research-spend-imbalance chapter as the smaller of the two main Australian funders of cardiac-arrest work.

NHMRC — National Health and Medical Research Council. Australia's main federal funder of health and medical research; sets the national grant portfolio analysed in the research-spend-imbalance chapter.

NSW Ambulance — The statutory emergency ambulance service for New South Wales, a division of NSW Health. Operator of the computer-aided dispatch system, the VACAR-aligned registry feed for NSW, and the GoodSAM integration.

VACAR — Victorian Ambulance Cardiac Arrest Registry. The long-running Victorian OHCA registry whose published survival trends and epidemiological denominators are used as the closest Australian benchmark for this report [11].

REFERENCES

References

APPENDIX M

Appendix M · Methodology: data-layer hierarchy and modelled demand allocation scoring

This chapter documents the data-layer hierarchy the LIFT model uses. Three distinct objects sit in the pipeline; they are not interchangeable, and conflating them was the single highest-risk presentational defect the Codex 5.4 adversarial review identified (RTV-001, 2026-04-23). Sprint 44 introduced the observed layers above the retained synthetic proxy; Sprint 46 names all three objects explicitly and relabels the SA2 surface as a modelled allocation score rather than an incidence-per-100,000 map.

The Greater Sydney observed envelope

The published NSW Ambulance Cardiac Arrest Registry reports approximately 6,000 to 6,700 attended OHCA events per year across NSW between 2017 and 2023, with the Greater Sydney footprint accounting for roughly 4,800 of those events per annum once the metropolitan denominator is allowed for [3,39]. Registry extracts at SA2 or SA3 granularity remain behind a restricted-data request (draft letter at `logs/letters/draft/nsw_ambulance_ohca_registry_data_request.md`); the state-aggregate annual outcomes table is committed as an observed fixture at `data/observed/ohca_registry/nsw_annual_outcomes.csv`. This envelope anchors the scaling total for the modelled allocation score below; it is not an incidence map.

The SA3 observed ambulance-demand proxy

BHI Healthcare Quarterly publishes quarterly ambulance activity and Category 1 response-time data at SA3 (SA4 where SA3 counts are suppressed) from 2010 onward, sourced from NSW Ambulance CAD (source register at `docs/sources/mapping_data_and_tools.md`; exact release citation will land in the Sprint 48 source-audit table). The Sprint 44 ingest at `data/observed/bhi/ambulance_activity_sa3.csv` carries 32 NSW SA3s spanning all three ABS Remoteness Area tiers. This layer is a demand-and-response proxy, not OHCA-specific; it is the observed basis for the three-tier candidate ranking at `outputs/tables/tier_candidate_sites.csv`. Category 1 response-time deficit against a tier-specific target (8 min Metro / 10 min Regional / 15 min Rural) is computed per SA3 and weights the ranking composite.

The SA2 modelled demand allocation score

The SA2-granular surface historically labelled as a synthetic OHCA count or incidence figure is in fact a **modelled demand allocation score**, not an incidence map. It prioritises SA2s for drone-base placement on the basis of an explicit, reproducible formula (population × age-65+ share × IRSD weight, scaled to the Greater Sydney envelope above); it is a planning prioritisation, not a published epidemiological incidence. The output lives at `data/derived/synthetic_ohca_hotspots.csv`. Readers should treat it as a modelled allocation score (useful for identifying high-priority SA2s at fine granularity where the Registry has not yet released SA2 counts) and not as observed incidence. Relabelling the surface in figures and captions as an allocation score rather than as an incidence figure is the Sprint 46 action that closes RTV-001. The scoring formula itself is unchanged; only the interpretation and presentation are corrected.

Source hierarchy (observed > derived > modelled > synthetic). The retained SEIFA-IRSD-weighted proxy described below is modelled, not synthetic; it is derived from observed ABS inputs by an explicit formula anchored to the published OHCA envelope. Its `data_label` column carries the legacy `synthetic` tag for backward compatibility with downstream tooling, but the analytical status is modelled. It is no longer the default analytical layer in any public-facing output where an observed or derived reading is available, and wherever a map caption, legend, or tooltip names the SA2 surface, the word used is “allocation score” rather than “rate” or “incidence”.

Motivation

An earlier version of the synthetic cohort (Sprint 01 checkpoint, later retained through Sprints 16–20) ranked SA2s by an age-dominated composite proxy. That proxy ordered three wealthy northern-shore SA2s, Chatswood (East)–Artarmon, Hurstville, and Wahroonga–Warrawee, at the top of the list. Clinically, that ordering is implausible: OHCA incidence is driven by age *and* socio-economic disadvantage, not by age alone. Wealthy suburbs carry older age structures but markedly better cardiovascular health profiles (higher preventive-care uptake, lower smoking prevalence, better access to primary care) than the disadvantaged south-west corridor (Fairfield, Cabramatta, Liverpool, Bankstown, Mount Druitt, Campbelltown), where absolute OHCA counts should be higher. Sprint 26 replaces the old composite with an explicit SEIFA-weighted formula so the proxy reflects this.

Formula

For each SA2 i:

$$\text{OHCA}_i \propto \text{population}_i \times \text{age65+ share}_i \times \text{IRSD weight}_i$$

where the IRSD weight is a monotonically decreasing function of the ABS 2021 SEIFA Index of Relative Socio-economic Disadvantage (IRSD) decile [20]:

$$\text{IRSD weight}_i = \frac{11 - \text{IRSD decile}_i}{10}$$

Decile 1 (most disadvantaged) maps to a weight of 1.0; decile 10 (least disadvantaged) maps to 0.1. A linear decile weighting is chosen for v1 because it is interpretable, monotone, and does not require committing to a specific published effect-size estimate; the published OHCA-by-deprivation gradient literature reports relative risks in the 1.5×–3× range between most- and least-advantaged quintiles (Australian and international cohorts), and a linear decile weighting sits conservatively inside that envelope. A log-linear or quintile-step weighting can be substituted when the NSW Ambulance cohort is joined in and the empirical gradient can be measured directly.

The product population × age65+ share gives the expected count of at-risk person-years; the IRSD weight then modulates that by cardiovascular-risk profile. Formally the proxy is a multiplicative hazard on person-years, not an age-standardised rate; the latter requires the NSW Ambulance event-level data.

Scaling and base-rate anchor

The per-SA2 raw weight is turned into an integer event count by linear scaling:

$$\text{scale} = \frac{T}{\sum_i \text{pop}_i \times \text{age65+}_i \times \text{IRSD}_i} \qquad \text{OHCA}_i = \text{round}\left(\text{raw}_i \times \text{scale}\right)$$

The target total $T = 4\,200$ cases per year is chosen so the sum across the 37 project SA2s lands inside a [3 500, 5 000] envelope consistent with:

- **Hasselqvist-Ax et al., 2015** [7]: the Swedish national registry paper reports a crude OHCA incidence of approximately 100 cases per 100 000 person-years, the closest well-characterised incidence we trust without a bespoke Australian fit.
- **NSW Ambulance Annual Report 2022–23** [39]: reports roughly 4 800 adult-attended OHCA events per year across the Greater Sydney footprint. The project 37-SA2 subset covers a population substantially smaller than the full metro, so a 4 200 target is the upper-mid point of the envelope after the ABS 2021 Greater Sydney denominator [10] is allowed for.

Tessa’s Sprint 26 test `test_csv_total_count_is_within_nsw_ambulance_envelope` fails the build if the emitted total falls outside the [3 500, 5 000] window; the invariant is enforced, not just documented.

Sensitivity to the SEIFA weighting choice

The Sprint 26 formula fixes a single linear decile weighting (11−decile)/10. A sophisticated reviewer is entitled to ask whether the ranking result is an artefact of that specific modelling choice. This subsection exercises two alternative weighting schemes end-to-end against the same curated inputs in `data/curated/sa2_greater_sydney.csv` and re-emits the synthetic count per SA2; the pass is documented in `scripts/citations/sensitivity/seifa_sensitivity.py` and is purely arithmetical; no new data sources enter. All outputs remain [SYNTHETIC] and pending verification against NSW Ambulance data access.

Variant A: log-linear (exponential) decile weighting. Substitute $\text{IRSD weight}_i = 2^{(11-\text{IRSD decile}_i)/10}$. This pulls the middle deciles up relative to the linear scheme and therefore gives mid-decile SA2s more weight.

Variant B: quintile-step weighting. Collapse the decile into a quintile-step function: deciles 1–2 → 1.00, 3–4 → 0.75, 5–6 → 0.55, 7–8 → 0.35, 9–10 → 0.20. This flattens the within-quintile gradient and therefore slightly down-weights the 1–2 extreme relative to the linear scheme.

Top-5 synthetic OHCA hotspots under the Sprint 26 linear baseline and the two sensitivity variants. Counts are rescaled to the same T = 4 200 target per variant so the columns are directly comparable. [SYNTHETIC] {#tbl:methodology-sensitivity}

Rank	Baseline linear $(11-d)/10$	Variant A: log-linear $2^{(11-d)/10}$	Variant B: quintile-step
1	Fairfield (d=1, 232)	Bankstown (d=2, 181)	Bankstown (d=2, 250)
2	Bankstown (d=2, 229)	Fairfield (d=1, 176)	Fairfield (d=1, 227)
3	Cabramatta – Lansvale (d=1, 221)	Cabramatta – Lansvale (d=1, 168)	Cabramatta – Lansvale (d=1, 216)
4	Campbelltown – South (d=1, 180)	Hurstville (d=5, 161)	Campbelltown – North (d=2, 188)
5	Campbelltown – North (d=2, 173)	Campbelltown – North (d=2, 136)	Granville – Clyde (d=2, 186)

Read the table by column, not by row. The **top-3** (Fairfield, Bankstown, Cabramatta – Lansvale) is stable across all three weighting schemes; only the internal ordering within the top-3 rotates (Bankstown and Fairfield swap rank 1 under the log-linear and quintile-step variants because Bankstown has a larger age-65+ population and the gap between decile-1 and decile-2 weight narrows under both alternatives). Positions 4 and 5 are less stable: the linear baseline places both Campbelltown SA2s in the top-5, the quintile-step variant drops Campbelltown – South to rank 6 in favour of Granville – Clyde (also decile 2), and the log-linear variant brings Hurstville (decile 5) into rank 4 because the exponential weight narrows the decile-1-to-decile-5 gap meaningfully. The 6th–20th band under all three variants is materially different between schemes and would not support a robust claim in prose.

Interpretation. The SEIFA-weighted proxy is robust to reasonable modelling choices at the top-3 of the ranking: the Fairfield / Bankstown / Cabramatta band sits at the head of the list under linear, log-linear, and quintile-step weightings alike; the 6th–20th band is more sensitive to the weighting choice and would benefit from NSW Ambulance cohort validation before any operational commitment is made downstream of it. A third variant, IRSAD substitution for IRSD, is not run here because the curated SA2 table carries only the IRSD decile; re-fetching the ABS SEIFA IRSAD datapack is the appropriate next sensitivity pass and is tracked as a deferred improvement for the sprint that onboards the NSW Ambulance incident data.

Inputs and provenance

Input	Source	Provenance label	Join key
Population count (per SA2)	ABS 2021 Census, G01 [10]	`observed+estimated`	`sa2_code`
Age 65+ share (%)	ABS 2021 Census, G04 [10]	`observed+estimated`	`sa2_code`
SEIFA IRSD decile	ABS 2021 SEIFA datapack, IRSD by SA2 [20]	`observed+estimated`	`sa2_code`
Centroid lat/lon	ABS 2021 ASGS SA2 digital boundaries (centroid)	`observed+estimated`	`sa2_code`

All inputs live in `data/curated/sa2_greater_sydney.csv` and are labelled `observed+estimated` in that file’s `data_label` column. The “estimated” component reflects that the SEIFA decile values carried in the curated CSV are the project’s best reconciliation of the ABS IRSD SA2 datapack pending a full re-fetch; values flagged `pending-verification` in the curated table would cascade into the synthetic CSV as zero-weighted rows rather than hand-filled numbers.

Top-5 SA2s after re-ranking

The re-ranked top of the table shifts decisively into Sydney’s south-west, the pattern Roger anticipated in the Sprint 26 brief when he observed that “there are lots of wealthy boomers in the northern suburbs, but unhealthy people in the west / south west.”

Top-5 synthetic OHCA hotspots after SEIFA IRSD weighting, with the input values the formula consumed for each SA2. Population and age share are ABS 2021 Census G01/G04 [10]; IRSD decile is ABS 2021 SEIFA [20]. {#tbl:methodology-top5}

Rank	SA2	Population	Age 65+ share	IRSD decile	Synthetic OHCA/year
1	Fairfield	13,800	15.2%	1	232
2	Bankstown	14,600	15.8%	2	229
3	Cabramatta - Lansvale	12,100	16.5%	1	221
4	Campbelltown - South	11,200	14.5%	1	180
5	Campbelltown - North	12,600	13.8%	2	173

All five SA2s sit in IRSD decile 1 or 2 (most disadvantaged). Wahroonga–Warrawee (IRSD decile 9) falls from rank 3 to rank 33, Chatswood (East)–Artarmon (decile 8) falls from rank 1 to outside the top-10, and the absolute counts redistribute mass toward the south-west, as visible in the Chapter 3 proportional-symbol map once it is regenerated from this CSV.

Limitations and pending verification

1. **Ecological fallacy.** A monotone IRSD weight captures the average deprivation gradient across an SA2 but cannot resolve within-SA2 variation. Empirical Bayes smoothing at SA1 granularity is the right refinement, and is deferred to the sprint that consumes the NSW Ambulance incident data.
2. **Age mid-point only.** The proxy uses the 65+ share as a single breakpoint. The real OHCA age curve is steeper in the 75+ band; once the NSW Ambulance data is joined, per-year age bins should replace the 65+ share.
3. **No sex stratification.** Male OHCA incidence is roughly 1.7× female in the comparable age bands [7], but the ABS SA2 sex split is close to 50/50 in the 37 project SA2s so the aggregate ranking is insensitive to this. The model should still be re-fit once a stratified NSW Ambulance series is available.
4. **Linear decile weighting is a placeholder.** The v1 linear weighting is deliberately conservative and has no empirical fit. When the NSW Ambulance cohort lands, re-fit *f*(IRSD) directly on the joined dataset.

5. **Output remains a proxy.** Every downstream figure, table, and shortlist entry must carry the “synthetic cohort, pending NSW Ambulance data access” caveat until the real feed replaces this CSV. The Chapter 3 spatial-analysis caveat paragraph points readers back to this chapter.

Reproduction

The CSV is regenerable from the curated inputs with a single-file Python routine:

```
count_i = round(
    population_i * (age_65_plus_pct_i / 100) *
    ((11 - irsd_decile_i) / 10) * scale
)
```

where `scale` is calibrated so the integer totals sum to the target $T = 4\,200$. Tessa’s Sprint 26 test `test_counts_are_consistent_with_formula` inverts this: it re-derives the scale from the dataset’s own totals and fails the build if any row deviates from formula prediction by more than one count.

Image Credits

This appendix records every image that appears in LIFT Study 01, its licence, and its source. It is generated from `content/assets/images/manifest/image-manifest.json` by `scripts/build-image-credits.js`; manual edits here will be overwritten on the next build. Corrections flow through the manifest.

Every image also carries alt text in the manifest; that alt text is wired into the published HTML and paged.js PDF via Paige’s templates so that screen-reader users receive a specific, clinical-register description of each figure’s content.

Bespoke and commissioned imagery

- **Figure 1** (`cover-lift-study-01`): LIFT Study 01 cover image: an AED-equipped delivery drone over suburban Sydney at sunrise, illustrating the programme’s core feasibility question. Cover image generated by Roger Lawrence (Sora, 2026), drone reference imagery courtesy of Karolinska Institutet press kit. licence: all-rights-reserved source: Roger Lawrence (Sora), No Kill Switch Research Programme
- **Figure 2** (`uas-sky-reference`): Karolinska Institutet drone reference photograph, used as the foundation image for the AI-generated LIFT Study 01 cover. Photo courtesy of Karolinska Institutet press kit. licence: press-kit-use source: Karolinska Institutet drone press kit (supplied via Roger)

- **Figure 3** (`lift-hero-aed-drone-suburb-sunrise`): Synthetic editorial hero image: cargo drone carrying an AED over a quiet Australian suburb at sunrise. LIFT Study 01 / synthetic image licence: all-rights-reserved source: LIFT visual asset pack 2026-04-25 (Sora generation)
- **Figure 4** (`lift-aed-delivery-moment`): Synthetic illustration of an AED-carrying drone delivering to a residential location. LIFT Study 01 / synthetic image licence: all-rights-reserved source: LIFT visual asset pack 2026-04-25 (Sora generation)
- **Figure 5** (`lift-dispatch-operations-centre`): Synthetic illustration of a dispatch operations centre coordinating drone-delivered AED response. LIFT Study 01 / synthetic image licence: all-rights-reserved source: LIFT visual asset pack 2026-04-25 (Sora generation)
- **Figure 6** (`lift-drone-dock-launch-readiness`): Synthetic illustration of a compact drone dock with an AED-delivery drone in ready posture. LIFT Study 01 / synthetic image licence: all-rights-reserved source: LIFT visual asset pack 2026-04-25 (Sora generation)
- **Figure 7** (`lift-human-impact-aed-responder`): Synthetic illustration of a community responder carrying a drone-delivered AED case toward a residence. LIFT Study 01 / synthetic image licence: all-rights-reserved source: LIFT visual asset pack 2026-04-25 (Sora generation)

Manufacturer press-kit imagery

- **Figure 8** (`lift-drone-airframe-wingcopter`): Wingcopter 198 fixed-wing VTOL cargo platform — one of the candidate airframes evaluated. Wingcopter media assets licence: Wingcopter media-assets coverage-use terms source: Wingcopter media assets

Creative Commons imagery

- **Figure 9** (`sydney-basemap`): Sydney base map (CartoDB Positron, © OpenStreetMap contributors) — provides coastline + road-network context beneath the SA2 data overlays. Basemap © OpenStreetMap contributors, tiles © CartoDB (Positron). licence: CC-BY-3.0 (CartoDB) + ODbL / CC BY 2.0 (OpenStreetMap upstream) source: CartoDB Positron (raster tiles stitched from OpenStreetMap data)
- **Figure 10** (`lift-aed-closeup`): An automated external defibrillator (AED) of the kind a drone-delivered network would carry. Mikhail Nilov / Pexels licence: Pexels License source: Mikhail Nilov / Pexels
- **Figure 11** (`lift-public-aed-cabinet`): A publicly accessible AED cabinet — illustrative of the fixed-installation public-access model the LIFT report critiques. Lydia / Wikimedia Commons (CC BY 2.0) licence: CC BY 2.0 source: Lydia / Wikimedia Commons
- **Figure 12** (`lift-public-aed-on-wall`): A wall-mounted AED — illustrative of the public-access AED footprint the LIFT report contrasts with on-demand drone delivery. Kaustrins / Wikimedia Commons (CC BY-SA 4.0) licence: CC BY-SA 4.0 source: Kaustrins / Wikimedia Commons
- **Figure 13** (`lift-problem-australian-suburb-dawn`): Australian suburban context — the residential geography where the majority of OHCA events occur. Break Media / Pexels licence: Pexels License source: Break Media / Pexels

- **Figure 14** (`lift-problem-suburban-house`): A quiet residential street — the kind of location where the majority of OHCA events occur, beyond the reach of public-access AEDs. David Yu / Pexels licence: Pexels License source: David Yu / Pexels

Category: diagrams

- **Figure 15** (`chain-of-survival-lift`): Chain of survival — the sequence from collapse to advanced care, with the LIFT intervention zone highlighted. LIFT Study 01 / No Kill Switch (CC BY 4.0) licence: CC BY 4.0 source: LIFT visual asset pack 2026-04-25
- **Figure 16** (`four-minute-response-model`): Illustrative four-minute response model — detection, launch, flight, delivery, bystander AED use. LIFT Study 01 / No Kill Switch (CC BY 4.0) licence: CC BY 4.0 source: LIFT visual asset pack 2026-04-25
- **Figure 17** (`dispatch-to-delivery-workflow`): Dispatch-to-delivery workflow — the operational sequence from 000 call recognition through ambulance handover. LIFT Study 01 / No Kill Switch (CC BY 4.0) licence: CC BY 4.0 source: LIFT visual asset pack 2026-04-25
- **Figure 18** (`system-trust-governance-loop`): System-trust governance loop — the operational and regulatory cycle that supports continuous safety assurance. LIFT Study 01 / No Kill Switch (CC BY 4.0) licence: CC BY 4.0 source: LIFT visual asset pack 2026-04-25

---

2 additional placeholder assets exist in the source repository but are not published in this document. They are retained for design-system regression fixtures only.

Appendix generated from content/assets/images/manifest/image-manifest.json — do not edit by hand.

Appendix A: Implications for Research

Research gap

Out-of-hospital cardiac arrest remains one of the leading causes of death in developed nations, with survival critically dependent on time-to-defibrillation [5,7,21]. While the survival benefit of early defibrillation is well established, public-access AED programmes have limited reach into residential settings where the majority of events occur [3,36]. Recent pilot programmes in Sweden have demonstrated that drone-delivered AEDs can arrive before ambulances and be used successfully by bystanders [19,30], but no study has addressed the specific optimisation challenges of deploying such networks in Australian suburban environments.

The gap is threefold. First, existing coverage models for drone AED delivery use simplified geometric assumptions (circular catchments) that do not account for the built environment, tree canopy, controlled airspace corridors, or weather-dependent operational envelopes specific to cities like Sydney. Second, the interaction between drone-delivered AED availability and bystander willingness, ability, and time-to-apply has not been studied in Australian populations; the Swedish evidence, while promising, reflects a population with substantially higher bystander CPR rates and AED familiarity. Third, no cost-effectiveness framework exists for drone AED networks that accounts for the opportunity cost relative to alternative interventions such as dispatcher-assisted CPR, fixed AED densification, or community responder programmes.

This feasibility model provides a preliminary analytical foundation (synthetic demand surfaces, proxy-based coverage estimates, and first-pass survival impact scenarios) but explicitly identifies the limitations that warrant rigorous research investigation.

Candidate research questions

RQ1: Spatial optimisation under operational constraints. What is the optimal placement and configuration of a drone AED base network in suburban Sydney when accounting for controlled airspace exclusion zones, weather-dependent flight envelope reductions, fleet maintenance downtime, and variable demand distributions across time-of-day and season?

This question extends the current project’s greedy set-cover approach to incorporate real-world constraints that reduce effective coverage from nominal geometric reach. Methods would include mixed-integer programming with stochastic demand inputs, calibrated against actual OHCA registry locations (subject to data access through NSW Ambulance research partnerships).

RQ2: Bystander retrieval and application dynamics. What factors influence the time from drone AED arrival at a residential address to bystander retrieval and successful shock delivery, and how does this retrieval interval modify the net survival benefit?

The current model assumes a fixed 1-minute retrieval time, but actual bystander behaviour in Australian suburban settings is uncharacterised. This question would require field trials or high-fidelity simulation studies, potentially incorporating HCI methods for payload design, dispatcher-guided retrieval protocols, and community training interventions.

RQ3: Comparative cost-effectiveness. What is the cost-effectiveness of a drone AED network compared to alternative defibrillation access strategies (fixed AED densification, community responder apps, dispatcher-assisted CPR enhancement) in suburban settings with varying demographic and geographic characteristics?

This question requires a decision-analytic model (likely a Markov model or microsimulation) incorporating the survival functions validated in RQ1 and RQ2, local cost data from procurement processes, and health system outcome data from the NSW OHCA registry.

Proposed methodology

The research programme would proceed in three integrated phases:

Phase 1 (12 months): Spatial model development and validation. Obtain geocoded OHCA registry data through a formal research agreement with NSW Ambulance. Develop a constrained spatial optimisation model incorporating ABS SA2/mesh block geography, CASA airspace layers, Bureau of Meteorology wind and precipitation data, and fleet availability models. Validate coverage predictions against a small-scale field trial using drone flight paths in a designated test corridor.

Phase 2 (12 months): Bystander interaction study. Conduct a mixed-methods study combining structured observation of simulated AED drone retrievals at residential addresses (n=100+), semi-structured interviews with participants, and analysis of dispatcher-bystander communication recordings from international AED drone programmes (through data-sharing agreements with the Karolinska group). Develop a validated retrieval-time distribution for Australian suburban settings.

Phase 3 (12 months): Cost-effectiveness modelling. Build a decision-analytic model comparing the drone AED network against three alternative strategies using outputs from Phases 1 and 2. Conduct probabilistic sensitivity analysis across key parameters. Submit findings to a health technology assessment audience.

Significance and expected contribution

This research programme would produce three outputs with direct translation potential. First, a validated spatial optimisation framework that could be applied to any Australian metropolitan area considering drone AED deployment (tools and methodology, not just Sydney-specific recommendations). Second, the first empirical data on Australian bystander AED retrieval behaviour, which has implications beyond drone delivery for all public-access defibrillation programmes. Third, a comparative cost-effectiveness analysis that would inform health system investment decisions and potentially support MRFF or NHMRC grant applications for implementation research.

The work is positioned at the intersection of emergency medicine, operations research, human-computer interaction, and health economics. It would suit a candidate with a quantitative health services research background and access to spatial analysis and health economic modelling expertise. The feasibility model produced by the current project provides an immediate starting point and a structured evidence base for a PhD proposal.

Data classification: This appendix is an estimated research framing based on the project's feasibility findings and the evidence catalogue (see outputs/evidence/evidence_catalogue.md).

1. American Heart Association. 2020 Guidelines for Cardiopulmonary Resuscitation and Emergency Cardiovascular Care. Circulation 142, Supplement 2. 2020. Available at: https://cpr.heart.org/en/resuscitation-science/cpr-and-ecc-guidelines.
2. Nichol G, Huszti E, Birnbaum A, Mahoney B, Weisfeldt M, Travers A, et al. Cost-Effectiveness of Lay Responder Defibrillation for Out-of-Hospital Cardiac Arrest. Annals of Emergency Medicine. 2009;54(2):226–235.e2. doi: 10.1016/j.annemergmed.2009.01.021.
3. Australian Institute of Health and Welfare. Out-of-hospital cardiac arrest in Australia: annual report 2021. Canberra: Australian Institute of Health and Welfare. 2021. Available at: https://www.aihw.gov.au/.
4. Waalewijn RA, Vos R de, Tijssen JGP, Koster RW. Survival models for out-of-hospital cardiopulmonary resuscitation from the perspectives of the bystander, the first responder, and the paramedic. Resuscitation. 2001;51(2):113–22. doi: 10.1016/s0300-9572(01)00407-5.
5. Valenzuela TD, Roe DJ, Cretin S, Spaite DW, Larsen MP. Estimating Effectiveness of Cardiac Arrest Interventions. Circulation. 1997;96(10):3308–13. doi: 10.1161/01.cir.96.10.3308.
6. Chan PS, Krumholz HM, Nichol G, Nallamothu BK. Delayed Time to Defibrillation after In-Hospital Cardiac Arrest. New England Journal of Medicine. 2008;358(1):9–17. doi: 10.1056/nejmoa0706467.
7. Hasselqvist-Ax I, Riva G, Herlitz J, Rosenqvist M, Hollenberg J, Nordberg P, et al. Early Cardiopulmonary Resuscitation in Out-of-Hospital Cardiac Arrest. New England Journal of Medicine. 2015;372(24):2307–15. doi: 10.1056/nejmoa1405796.
8. National Rural Health Alliance. Heart of the Nation: A movement to save lives. Partyline (National Rural Health Alliance). 2024. Available at: https://www.ruralhealth.org.au/partyline/article/heart-nation-movement-save-lives.html.
9. Central Western Daily. Stolen public defibrillator found in Orange bush by NSW Police. Central Western Daily. 2024. Available at: https://www.centralwesterndaily.com.au/story/8673282/stolen-public-defibrillator-found-in-orange-bush-by-nsw-police/.

10. Australian Bureau of Statistics. 2021 Census of Population and Housing: Greater Sydney SA2 profiles. Canberra: Australian Bureau of Statistics. 2022. Available at: https://www.abs.gov.au/census.

11. Nehme E, Anderson D, Salathiel R, Carlyon A, Stub D, Cameron PA, et al. Out-of-hospital cardiac arrests in Victoria, 2003–2022: retrospective analysis of Victorian Ambulance Cardiac Arrest Registry data. Medical Journal of Australia. 2024;221(11):603–11. doi: 10.5694/mja2.52532.

12. Morrison A, Arnold J. The NSW Ambulance GoodSAM program – a review of the first 12 months. Council of Ambulance Authorities Congress 2024 (conference presentation and abstract). 2024. Available at: https://caacongress.net.au/the-nsw-ambulance-goodsam-program-a-review-of-the-first-12-months/.

13. NSW Government. 100 lives saved as NSW Ambulance GoodSAM volunteers answer the call. Ministerial release, Minister for Health (NSW). 2026. Available at: https://www.nsw.gov.au/ministerial-releases/100-lives-saved-as-nsw-ambulance-goodsam-volunteers-answer-call.

14. Ringh M, Rosenqvist M, Hollenberg J, Jonsson M, Fredman D, Nordberg P, et al. Mobile-Phone Dispatch of Laypersons for CPR in Out-of-Hospital Cardiac Arrest. New England Journal of Medicine. 2015;372(24):2316–25. doi: 10.1056/nejmoa1406038.

15. Smith CM, Lall R, Fothergill RT, Spaight R, Perkins GD. The effect of the GoodSAM volunteer first-responder app on survival to hospital discharge following out-of-hospital cardiac arrest. European Heart Journal Acute Cardiovascular Care. 2022;11(1):20–31. doi: 10.1093/ehjacc/zuab103.

16. Delardes B, Gregers MCT, Nehme E, Ray M, Hall D, Walker T, et al. Smartphone-activated volunteer responders and survival to discharge after out-of-hospital cardiac arrests in Victoria, 2018–23: an observational cohort study. Medical Journal of Australia. 2025;222(10):504–9. doi: 10.5694/mja2.52673.

17. Lijovic M, Bernard S, Nehme Z, Walker T, Smith K. Public access defibrillation—Results from the Victorian Ambulance Cardiac Arrest Registry. Resuscitation. 2014;85(12):1739–44. doi: 10.1016/j.resuscitation.2014.10.005.

18. Schierbeck S, Nord A, Svensson L, Ringh M, Nordberg P, Hollenberg J, et al. Drone delivery of automated external defibrillators compared with ambulance arrival in real-life suspected out-of-hospital cardiac arrests: a prospective observational study in Sweden. The Lancet Digital Health. 2023;5(12):e862–71. doi: 10.1016/s2589-7500(23)00161-9.

19. Schierbeck S, Svensson L, Claesson A. Use of a Drone-Delivered Automated External Defibrillator in an Out-of-Hospital Cardiac Arrest. New England Journal of Medicine. 2022;386(20):1953–4. doi: 10.1056/nejmc2200833.

20. Australian Bureau of Statistics. Socio-Economic Indexes for Areas (SEIFA), Australia, 2021: Index of Relative Socio-economic Disadvantage (IRSD) by Statistical Area Level 2. Canberra: Australian Bureau of Statistics. 2023. Available at:

https://www.abs.gov.au/statistics/people/people-and-communities/socio-economic-indexes-areas-seifa-australia/2021.

21. Larsen MP, Eisenberg MS, Cummins RO, Hallstrom AP. Predicting survival from out-of-hospital cardiac arrest: A graphic model. Annals of Emergency Medicine. 1993;22(11):1652–8. doi: 10.1016/s0196-0644(05)81302-2.

22. Swoop Aero. Healthcare logistics in Malawi and the Pacific. Company website. 2024. Available at: https://www.swoop.aero/.

23. Civil Aviation Safety Authority. Advisory Circular 101-01: Remotely piloted aircraft systems — licensing and operations. Civil Aviation Safety Authority (CASA), Canberra. 2024. Available at: https://www.casa.gov.au/sites/default/files/2021-08/advisory-circular-101-01-remotely-piloted-aircraft-systems-licensing-and-operations.pdf.

24. Civil Aviation Safety Authority. Remote pilot operator certificates (ReOC): commercial and beyond-visual-line-of-sight operations. Civil Aviation Safety Authority (CASA), Canberra. 2024. Available at: https://www.casa.gov.au/drones/drone-rules/reoc.

25. Civil Aviation Safety Authority. Part 101 Manual of Standards and Advisory Circulars: unmanned aircraft and rockets. Canberra: Civil Aviation Safety Authority. 2024. Available at: https://www.casa.gov.au/.

26. Wing Aviation. Wing Logan, Queensland: commercial delivery operations and autonomous BVLOS flights. Wing Aviation company website. 2023. Available at: https://wing.com/australia/logan/.

27. Manna Aero. Manna drone delivery operations: autonomous last-mile flights in Dublin, Ireland. Manna Aero company website. 2023. Available at: https://www.manna.aero/.

28. Zipline International. Zipline Rwanda operations: sustained medical drone logistics. Company website. 2024. Available at: https://www.flyzipline.com/rwanda.

29. Scalea JR, Restaino S, Scassero M, Blankenship G, Bartlett ST, Wereley N. An initial investigation of unmanned aircraft systems (UAS) and real-time organ status measurement for transporting human organs. IEEE Journal of Translational Engineering in Health and Medicine. 2018;6:1–7. doi: 10.1109/JTEHM.2018.2875704.

30. Everdrone AB. First-ever delivery of defibrillator drone saves a life. Company website. 2022. Available at: https://www.everdrone.com/first-ever-delivery-of-defibrillator-drone-to-save-life/.

31. Coute RA, Panchal AR, Mader TJ, Neumar RW. National Institutes of Health–Funded Cardiac Arrest Research: A 10-Year Trend Analysis. Journal of the American Heart Association. 2017;6(7). doi: 10.1161/jaha.116.005239.

32. Coute RA, Panchal AR, Mader TJ, Kurz MC. Research Funding of the Top Science Gaps in the American Heart Association Cardiac Arrest Guidelines. Circulation: Cardiovascular Quality and Outcomes. 2021;14(5). doi: 10.1161/CIRCOUTCOMES.120.007627.

33. Coute RA, Mader TJ, Kurz MC. Evaluation of National Institutes of Health cardiac arrest research based on “chain of survival” links. Academic Emergency Medicine. 2022;29(11). doi: 10.1111/acem.14569.

34. Coute RA, Huebinger R, Perman SM, Del Rios M, Kurz MC. Evaluating the National Institutes of Health Pipeline for Resuscitation Science Investigators. Journal of the American Heart Association. 2024;13(12). doi: 10.1161/JAHA.124.035854.

35. Pointon S, Ingles J, Timbs S, Vandenberg J, Bray J, Page G, et al. National Funding for Cardiac Arrest Research in Australia: An Analysis of Competitive Grant Schemes. Heart, Lung and Circulation. 2026;35(4). doi: 10.1016/j.hlc.2025.11.013.

36. Deakin CD, Anfield S, Hodgetts GA. Underutilisation of public access defibrillation is related to retrieval distance and time-dependent availability. Heart. 2018;104(16):1339–43. doi: 10.1136/heartjnl-2018-312998.

37. Medicine I of. Strategies to Improve Cardiac Arrest Survival. Washington, DC: National Academies Press. 2015. doi: 10.17226/21723.

38. Callaway CW. Increasing National Institutes of Health Funding for Cardiac Arrest Research. Journal of the American Heart Association. 2017;6(7). doi: 10.1161/JAHA.117.006739.

39. NSW Ambulance. Annual Report 2022-23. Sydney: NSW Ministry of Health. 2023. Available at: https://www.ambulance.nsw.gov.au/.