TECHNICAL COMPANION

Joint Architecture for Forward Denial

Operational Concept and Technical Description

The Joint Architecture is grounded in the transparent-battlefield lessons of Ukraine and the Gulf, the operational geography of Australia's northern approaches, and the technological context of 2026. It is built to meet the requirement the 2026 National Defence Strategy sets: to deter through denial any attempt to project power against Australia through those approaches.

This document is the companion to the briefs 'The Changing Maritime Landscape' and the 'Joint Architecture for Forward Denial'.

Malcolm Edgar

Edited by Stephen Hayes MBE

August 2026

Version 1.0

In memory of the Australians and Americans
who fought the Battle of the Bismarck Sea, 2–4 March 1943.
They held the approaches together.

Provenance and sources. *This volume is derived entirely from open sources. Every factual claim is drawn from the publicly available materials cited in the reference lists (official government publications, parliamentary and congressional reporting, institutional research, manufacturer disclosures, and reputable press reporting), or is explicitly identified in the text as an estimate or analytic judgment constructed from those materials. No classified, protectively marked, or otherwise privileged information has been used in or has contributed to this document, and nothing in it draws upon, confirms, denies, or should be taken to imply any information held by any government beyond what the cited public sources themselves establish. Where the public record is limited, figures are flagged as planning estimates and their precision should be read accordingly. The requirements, design figures and effectiveness modelling are engineering estimates offered to structure the analysis rather than measurements, and each is stated with its basis so that a reader can test it.*

About the author. *Malcolm Edgar is a retired software architect and chief technology officer. Across a thirty-year career he was Chief Technology Officer of the fintech company Avoka Technologies, a Vice President of the Apache Software Foundation and founder of one of its open-source projects, and holds a patent in encryption methods. His career began at ADI (later Thales Australia, 1995–2000), building joint command-support systems for Headquarters ADF under JP2030 and maritime command-support software for the RAN frigate program. That is his only professional connection to Defence, disclosed here. He writes as a systems architect, not a naval professional: this volume's method (single-responsibility components, graceful degradation, no single point of failure) is the discipline of large-scale distributed systems applied to force design.*

About the editor. *Stephen Hayes MBE served more than thirty years in the Royal Australian Air Force and was the founding Chief Executive of the International Centre for Complex Project Management. He has served as the National Defence Industry Workforce and Skills Facilitator, appointed by the Minister for Defence Industry, and as Secretary of the Australian Naval Shipbuilding Industry Group to December 2025. He chairs the Gravity Group and the Global Access Partners Standing Committee on National Resilience.*

Contents

Glossary

The architecture	
Battery	Sixteen 40-foot high-cube launch containers at nine rounds each: 144 rounds.
JSTN	Joint Sovereign Targeting Network. Space-based SAR and emitter geolocation fused with JORN and the air picture: the network that finds targets and holds them for the fires, under Australian authority over tasking and release.
LRPF	Long-Range Precision Fires. The architecture's land-based fires, fielded as batteries of container-launched rounds.
LRSM	Long-Range Strike Missile. A cruise missile class munition with 2,400 km reach, 250 kg warhead, terminal sea-skim, passive machine-vision seeker, about US$300,000.
Sensing, custody and command	
C2	Command and Control.
Custody	Continuous knowledge of where a force is, current enough to guide a weapon.
ELINT	Electronic intelligence. Detection, identification and geolocation of an emitter (a warship's radars identify its class).
ELINT/datalink shell	The upper of the JSTN's two low-earth-orbit constellations: of the order of fifty to sixty-six cross-linked satellites at ~780 km carrying passive ELINT receivers and the assured L-band bearer that carries cueing, tactical data link and mid-course updates to the missiles in flight.
EO / IR	Electro-Optical and Infrared imaging. Daylight- and weather-limited, and drawn from commercial sources rather than built.
GNSS	Satellite navigation. Assumed to be denied in the engagements this volume describes.
JORN	Jindalee Operational Radar Network. Australia's over-the-horizon radar (OTHR), holding wide-area watch across the approaches.
SAR	Synthetic-Aperture Radar. Imaging radar that sees through cloud and darkness, by day or night.
SAR shell	The lower of the JSTN's two low-earth-orbit constellations: of the order of sixty small satellites at ~520–580 km carrying synthetic-aperture radar, giving all-weather detection and revisit across the approaches. Purchased from an allied supplier, Australian-owned and Australian-operated.
Sovereign core	The JSTN's integrating software, and the element that must be Australian-held outright: multi-sensor fusion into a single joint picture, command and control of engagements across dispersed shooters, and the weapon datalink that carries mid-course updates to the missiles in flight (Section 2.6).
Threat and engagement	
Exhaustion	The point at which a defended group has spent its entire interceptor magazine.
Mission-kill	Rendering a combatant unable to fight or fly, without sinking it. Destruction goes further, ending the ship as a fighting system. Sinking a large hull is neither claimed nor required: the deterrent is achieved at mission-kill.
SAG	Surface Action Group: surface combatants with a replenishment ship. A carrier group is the same problem with an air wing added.
SAM / VLS	Surface-to-air missile, carried in vertical-launch system cells. VLS cells cannot be reloaded at sea, the asymmetry the magazine argument rests on.
SEAD	Suppression of Enemy Air Defences: opening a corridor through which crewed aircraft fly.

Australian policy and programs	
GOCO	Government-Owned, Contractor-Operated: the state owns the plant, design authority and intellectual property; a commercial operator runs production. With vertical integration, additive manufacturing and a design authority close to the line, this is the procurement-model logic the cost figures rest on. The same systems bought cost-plus through a single prime price an order of magnitude higher.
GWEO	Guided Weapons and Explosive Ordnance enterprise: Australia's domestic munitions program. It supplies the fires' industrial setting and components; the sovereign missile line stands up alongside it (Sections 5.3–5.4).
IAMD	Integrated Air and Missile Defence: the layered shield over the bases and the force.
NDS / IIP	National Defence Strategy (2024, 2026) and Integrated Investment Program: the policy and the funded program it directs.
Self-reliance	Capability Australia can operate, control and resupply on its own, built with and through partners, where feasible and advisable (the sense the 2026 NDS uses). Sovereign, where it appears, means only decision and release authority, design authority and intellectual property held in Australia, or assured resupply from an Australian line. Alliances are force multipliers that widen the Government's choices.

Section 1 — Joint Architecture Operational Concept

This section states what the Joint Architecture is, then shows it working across five scenarios on the escalation ladder from peacetime to conflict. Each is illustrative, not predictive; and stays within the declared strategy of denial.

1.1 Strategic objective

The 2026 National Defence Strategy directs the Australian Defence Force to "**deter through denial any potential adversary's attempt to project power against Australia through our northern approaches**." Denial requires the standing ability to see a power-projection force in those approaches and to hold it at risk there, at the range the approaches impose, some thousands of kilometres from the mainland, across the archipelago and the sea-air gap.

That ability has two parts. The first is targeting-grade custody: a continuous, all-weather track of what is moving in the approaches, held to a quality that can direct a strike rather than merely report a contact. The second is long-range precision fires in enough depth to make a projection attempt costly across the whole of the approaches, not at a single point. This is the capability the architecture is built to supply.

It meets the objective on self-reliant terms. Custody and fires are generated from dispersed Australian sites and reconstituted from a domestic production line, so the capability holds whether or not allied support is available in a crisis (the test the Strategy of Denial sets). It expands the Government's choices rather than narrowing them, and works as a force multiplier alongside the integrated force and Australia's alliances: a near-term, additional layer of denial, not a replacement for any existing programme.

The sensing that enables denial also has a cooperative use. The same wide-area, all-weather surveillance that holds the approaches under custody detects unregistered and "dark" vessels across the maritime domain, the awareness Pacific partners need to police their own exclusive economic zones. The Strategy calls for "increasing coordination with regional partners to strengthen and preserve the global rules-based order"; maritime awareness Australia can share to that end is a direct contribution, and a standing sign that the architecture strengthens the region as well as defending the approaches.

1.2 Alignment with the Integrated Investment Program

The 2026 Integrated Investment Program commits $425 billion over the decade ($53 billion of it new) and names seven capability areas for accelerated delivery. This architecture advances four of them from a single system:

- **More lethal maritime capabilities**. Holding a maritime force at risk across the northern approaches is the architecture's driving role: a land-based magazine that adds lethal reach against ships and complements the fleet.
- **Long-range strike**. The Program's direction to "further expand long-range strike capabilities across the integrated, focused force" is delivered directly: heavy precision fires to 2,400 km, in depth.

- **Resilient multi-orbit space**. The commitment to secure, multi-orbit space capability that the Program funds for communications is the same posture the architecture's space-based custody of the approaches depends on.
- **Autonomous and collaborative fires**. The missiles are autonomous and networked: fires cued by any sensor on one shared targeting picture.

1.3 Capability overview

Stated as explicit capabilities, the architecture delivers:

Sensing and custody (JSTN)

- **Persistent space-based RF/ELINT** detection and geolocation of emitters (surface-search and navigation radars, and VHF/UHF communications, across roughly 150 MHz–15 GHz) at global coverage.
- **Non-cooperative ("dark") vessel detection** down to ~25 m length regardless of emissions, sweeping the ~69 million km² area of interest (~13% of the Earth's surface) every 5–10 hours.
- On designation, **custody-grade SAR revisit of ≤15 minutes**, tightening to minutes on cue.
- **Spotlight and dwell SAR imaging down to 25 cm** resolution.

Fires and reach (LRPF)

- **Heavy precision strike to ~2,400 km,** covering the same **~69 million km² / ~13% of the Earth's surface** from dispersed Australian sites: sensing and fires share one perimeter.
- **3,600 rounds in 25 batteries (20 forward, 5 reserve)**; each battery is sixteen 40-foot containers at nine rounds, 144 rounds.
- **250 kg blast-fragmentation warhead, three fuze settings:** point/impact, delayed (penetrating), and airburst (height-of-burst).
- **~900 tonnes of warhead** across the full magazine (3,600 × 250 kg).

System qualities

- **Network-centric**. The architecture is a kill web, not a single kill chain: any sensor can cue any shooter, so a track held by a satellite, JORN, a ship or an aircraft is serviced by whichever shooter is best placed. Collaborative Army–RAAF–Navy engagements become routine: land fires massing a synchronised time-on-target. With the magazine already forward and cued from space, fires stay responsive around the clock, and redundant sensor-to-shooter paths mean no single counterstrike breaks the chain.
- **Manufacture at rate.** An Australian government-owned line can build and reconstitute the full 3,600-round magazine within a year. The deterrent rests on that production rate rather than on the stockpile.
- **Low manning.** The batteries are containerised, statically emplaced and remotely commanded.

Figure 1.1 *The Joint Architecture: coverage. Sovereign sensing custody (satellite ground tracks and JORN) with the 2,400 km fires perimeter over the approaches. Battery sites illustrative; the force distributes 20 batteries across five sectors, with 5 held in reserve.*

LRPF battery = 16 containers = 144 rounds · ~US$43M

Sixteen identical 40-foot high-cube containers. Each holds nine rounds, three across by three high.

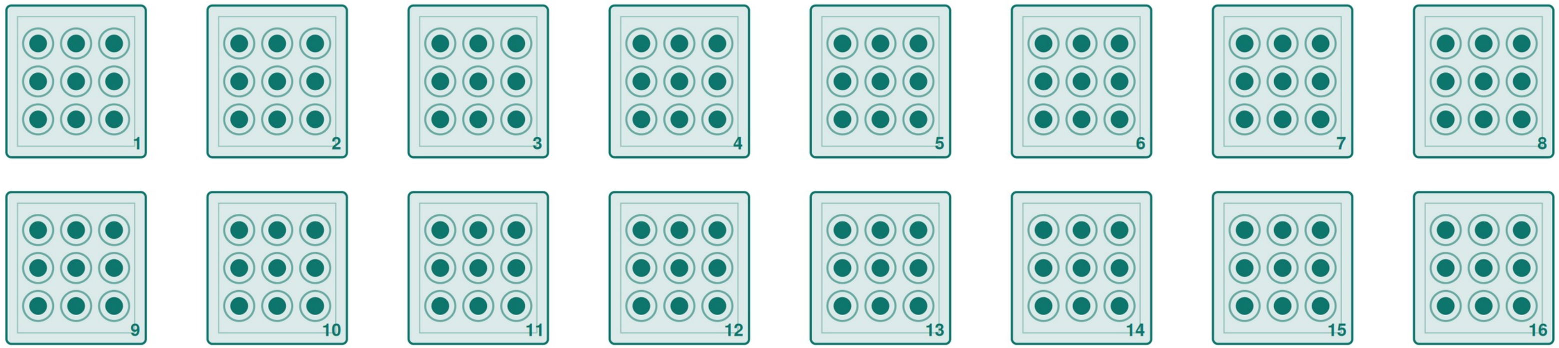

16 containers · 9 rounds each · 144 rounds, every one a ship-killer

THE ROUND

Reach	2,400 km
Length	6.75 m including booster
Body diameter	0.58 m
All-up mass	~1,650 kg, including ~800 kg of fuel
Propulsion	turbojet, additively manufactured
Profile	subsonic cruise; terminal sea-skim at 10 to 15 m
Warhead	250 kg blast-fragmentation, point, delayed fuze or air burst
Navigation	GNSS and inertial, JSTN mid-course update
Terminal guidance	passive machine-vision seeker; anti-radiation variant (SEAD)
Unit cost	~US$300,000
Per container	9 rounds, three across by three high
Container gross mass	~23 t against an ISO limit of 30.5 t

Anti-ship · land strike · suppression of air defences · anti-materiel

WHAT A BATTERY DOES

1 battery — **Small surface action group**
1 cruiser, 1 frigate, oiler · 124 rounds to destroy

2 batteries — **Large surface action group**
2 cruisers, 2 frigates, oiler · 235 rounds to destroy

3 batteries — **Carrier group**
carrier, 2 cruisers, 3 frigates · 346 rounds to destroy

THE FORCE

25 batteries · 400 containers · 3,600 rounds

~US$1.08 billion fly-away · 20 forward across five sectors, 5 in reserve

Figure 1.2 *The LRPF battery: sixteen 40-foot high-cube containers, nine rounds each, 144 rounds, about US$43 million. The missile, its seeker options including the anti-radiation variant for suppression of air defences, and what a battery does against each threat case. Figures are illustrative planning estimates; the derivations are at Sections 4 and 5.*

1.4 Scenario one — Task Group 107 custody

February 2025, as it happened. A People's Liberation Army Navy (PLAN) task group (a Renhai-class cruiser, a frigate and a replenishment ship) deployed through the Coral Sea and into the Tasman Sea, where it conducted live-fire exercises. First notice of the firings reached Australia from a commercial airline pilot who heard the broadcast warning in flight. Holding custody of three ships then consumed a national surge of warships and maritime patrol aircraft for weeks, as the group circumnavigated the continent. The deployment was lawful; the response was improvised.

The same deployment, with the network up. The group is acquired by wide-area radar surveillance as it enters the approaches and is in continuous custody from that hour: position, course and speed refreshed every few minutes by the SAR shell, emitters identified and fixed by the ELINT/datalink shell each time a radar or datalink transmits, the fused track carried in the joint picture alongside JORN and Wedgetail. The exercise broadcast is verified by observation rather than relayed by an airliner. No surge is required: the ships and aircraft that shadowed the 2025 group remain on their programmed tasks, because custody is a standing product of a constellation that is watching anyway. The track is releasable at Australian discretion: to the public as a published track history, to New Zealand and other partners as a coalition contribution. The group's commander knows all of this, which is the point: a force that is visibly held in custody from arrival to departure has lost the first advantage of operating at range, and Australia has demonstrated a capability without a single sortie.

1.5 Scenario two — coercion ladder answered in rungs

The campaign. A surface action group (SAG) takes station in the approaches off the North West Shelf as visible pressure in a dispute with presence as coercion. Coercion requires the group to be seen, to loiter and to stay, and each of those requirements is a vulnerability against a standing architecture. What government holds is a ladder of responses in which every rung is reversible, every rung imposes cost, and none of the early ones kills anyone.

The rungs, in order. First, custody made explicit: a public statement of the group's position and track history, demonstrating that it is held. Second, the conventional response, and the one available today: a maritime patrol aircraft on task overhead, or a surface combatant sailed north. It signals attention and it is visible, but at that distance from the continent it is a statement rather than a threat. One aircraft, or one hull, against a group carrying its own air defence, and the group's commander can do that arithmetic as readily as anyone. Third, visible posture from the architecture itself: a reserve battery is moved forward into the contested sector, and the adversary's own reconnaissance confirms the move, which is the message. Fourth, and the rung that carries the scenario: a single round, fired on an announced track that crosses ahead of the group and terminates in empty ocean.

What one round says. It is a demonstration, not an engagement: no target is struck, nobody is hurt, and the act is deliberately legible. But it settles four questions at once that no statement could settle. The custody is fire-control grade, because the round arrived where the group was rather than where it had been. The reach from Australian soil is real. The missile works, because it flew the profile and terminated where it was told. And the magazine behind it is deep, because a state that can spend one round to make a point is not husbanding a scarce inventory. The group's commander can verify every one of those from his own sensors within the minute, and the arithmetic he does next is his own. The group withdraws. That is the outcome the ladder exists to produce: coercion answered well below the threshold of violence, by a capability whose credibility was demonstrated rather than asserted.

1.6 Scenario three — maritime strike with fires

The setting. The ladder of Scenario two has been climbed and has not worked: the coercion continues, the demonstration has been noted and discounted, and government has authorised engagement. The group is the composition observed in 2025: a cruiser, a frigate and a replenishment ship. The chain runs from custody to assessment as one closed loop, every link of it Australian.

Table 1.1 *Scenario three: the kill chain, phase by phase.*

Phase	What happens
Detect	The SAR shell holds the group in custody, as it has since the ships entered the approaches; revisit every few minutes.
Classify	The ELINT/datalink shell fixes and identifies each emitter; cross-cued SAR imaging confirms hull classes. The fused track is fire-control grade.
Decide	Targeting data is released to the batteries under Australian national authority. No ally is asked; no foreign release gate exists in the chain.
Launch	A salvo of some 126 rounds fires from dispersed containers across two sectors, on multiple axes, sequenced so that all arrive together.
Mid-course	Over a flight of one to three hours the group moves tens of kilometres. The ELINT/datalink shell carries updates to the missiles in flight; the weapon datalink is the network's, not a partner's.
Terminal	The missiles descend to sea-skimming height for the final approach, seekers emitting nothing. The defence must engage every inbound track as a ship-killer, because every one is.
Assess	The next SAR pass arrives within minutes of the last impact. Imaging shows the group composition and track motion; the ELINT/datalink shell reports the group's radars off the air: the emitters that fell silent are the assessment. No aircraft flies battle-damage assessment over a defended group.
Re-strike or stand down	With custody unbroken and assessment in hand, a second salvo from the sector's remaining rounds needs only a decision.

The expenditure is some US$38 million in rounds against a group whose replacement value is measured in billions. The force holds twenty-four more engagements of the same weight with a production line behind them. The numbers are the model's at Section 4.

1.7 Scenario four — joint maritime strike with hypersonics

What has changed. The chassis is the same but now the Royal Australian Air Force fields HACM, the air-launched hypersonic weapon arriving through AUKUS Pillar 2. No official range is published; open sources put it between five hundred and a thousand kilometres.

The network is what lets the aircraft shoot. At HACM's reach the launch aircraft cannot see its own target: the radar horizon from a fighter's cruising altitude is about four hundred and sixty kilometres against a warship's mast, and beyond that no altitude an aircraft can reach recovers the line of sight. Nor can it look for itself, because the group detects an emitter long before that emitter detects the group. The network supplies the track instead, so the F/A-18 launches from the outer edge of the weapon's envelope, on a bearing of its own choosing, without ever radiating. That is collaborative engagement in the fullest sense: one network holding the target, several shooters firing on it, and the same network reconciling their different launch points and times of flight into a single simultaneous time on target, so that hypersonic weapons from the air and long-range fires from the land arrive together against a defence that must solve both at once.

The engagement. The target is a larger group, or a carrier group with its escorts. The hypersonic weapons go for the eyes: the fixed phased-array faces on which the group's air defence depends, engaged at hypersonic speed with almost no useful reaction time. In the same seconds the LRPF salvo arrives in its terminal phase, across multiple axes at sea-skimming height. The defence is asked to solve two incompatible problems at once: a handful of tracks it cannot react to, and a saturating raid it must service one round at a time. By combining RAAF hypersonics to suppress the group's air defences and LRPF to destroy the group, expensive munition use can be optimised.

What the joint architecture does with new capabilities. By plugging in new capabilities a network-centric joint force becomes greater than the sum of its parts, with each new capability and platform providing additional leverage and warfighting modes. With the addition of SSN boats the JSTN can provide timely updates of adversary movements in the open oceans.

1.8 Scenario five — joint land attack with SEAD

The build-up, watched before it is opposed. An adversary establishes an expeditionary lodgement on an archipelagic feature within the approaches, some 2,000 km from the Australian mainland: an airstrip, a port, dispersed surface-to-air missile (SAM) launchers and the radars that serve them. The site is designated, so the SAR satellites images it in spotlight mode (sub-metre, down to about 25 cm) every few hours through cloud and darkness. Revetments appear, transporter-erector vehicles move, a surveillance radar is deployed, aircraft arrive and park, and each is resolvable as an individual object. SAR imagery is combined with allied or commercial EO imagery to provide additional assessment. Meanwhile the ELINT/datalink shell fixes emitters the moment they radiate and builds the electronic order of battle over days. Reconnaissance produces the target set: surveyed aimpoints for the fixed infrastructure, fixed emitter positions for the air-defence system, and a current count of what is parked where.

Scene matching, and what the survey provides. The missile finds its aimpoint by matching what its seeker sees against a reference scene carried on board. The reference is built in advance from whatever imagery is available and current (sovereign SAR, allied and commercial optical, and prior survey) and for a site like this one it comes from the reconnaissance described above rather than from a peacetime library, because the site did not exist in peacetime. Sub-metre spotlight imagery is what makes the aimpoint precise enough to matter; the other sources add confidence and currency to the scene the seeker will actually see.

Phase one: suppression, by the fires. Missiles of the anti-radiation seeker variant go against the radars that are emitting; missiles with the EO seeker go against surveyed aimpoints, with airburst fuzing over the launcher revetments and the dispersal pans where aircraft are parked. The two work against each other’s counter: a radar that keeps emitting is engaged by the anti-radiation missile, and a radar that goes dark to survive has stopped doing its job and leaves its launchers to be serviced by the surveyed ones. Launchers destroyed in their revetments and aircraft destroyed on the ground do not have to be fought in the air. A salvo of this scale (two surface-to-air missile batteries with their radars, launchers and point defence, and the aircraft parked in the open) is of the order of fifty rounds, about a third of one battery. All of it is delivered from containers on Australian soil, with no aircrew committed.

Phase two: the strike package, through the corridor. The scale of the RAAF package is worth noting. With the objective at some 2,000 km the mission is built around tanking: eight to twelve strike aircraft, EA-18G Growlers for electronic attack, an E-7A Wedgetail for battle management, and KC-30A tankers meeting them at refuelling brackets outbound and inbound. That is on the order of twenty aircraft and a large fraction of the deployable air combat force, airborne for six to eight hours. Effective SEAD enables the package to run in and release heavy satellite-guided bombs of the JDAM class: a fraction of the cost, several times the warhead, and enough weight to break hardened targets. In the opening and most dangerous minutes of the mission aircrews and platforms have not been put at risk.

Phase three: assessment, by the network. Assessment is a function of the same satellite shells that built the target set. Within minutes of the strike the SAR layer re-images the site and the ELINT layer reports which emitters have fallen silent and which have not, with allied and commercial optical imagery adding further details. There is no separate reconnaissance sortie.

Phase four: the follow-up, again by the fires. What survives is re-engaged by the fires rather than by regenerating the air package. A second salvo can be fired within the hour, against aimpoints the assessment has just refreshed, while the aircraft are recovering, the crews resting and the tankers cycling for a sortie that is available tomorrow rather than this afternoon.

This is economy of force. Priced across the whole mission (rounds fired, weapons dropped and hours flown) the joint operation is also modelled as the cheaper one, at roughly half the cost of servicing the same target set from the air alone, because suppression converts expensive stand-off weapons into cheap direct-attack ones.

Section 2 — Joint Sovereign Targeting Network: Requirements and Design Concept

The JSTN is the architecture's keystone: the joint air-and-maritime sensing-and-targeting network on which every shooter depends. It is therefore the part that most needs to be specified rather than asserted. This section sets out what the network must do, a design concept that meets it, and the sequence by which it is fielded. It is a concept, not an engineering specification. Its purpose is to show that the keystone is buildable, on a timeline that matters, out of things that already exist.

2.1 Purpose and the joint imperative

The JSTN exists to find, track, maintain custody of, and cue fires against air and surface forces threatening Australia's approaches: to targeting-grade fidelity, on Australian authority, and resiliently enough to survive an attempt to blind it.

It is joint by necessity, not by aspiration. A denial engagement interleaves domains. A surface action group arrives with airborne early warning overhead, strike aircraft in support and submarines beneath, while Australia must simultaneously hold an air picture over its own force and over the critical assets. The network that finds ships and the network that tracks aircraft are not separable systems. They are one fused targeting backbone feeding both the long-range precision fires and integrated air and missile defence. The JSTN is a joint-force capability, owned across the services, and it is the sensing half of the declared task of deterring force projection against us.

One scope boundary is deliberate. The JSTN's remit is the surface and air picture, together with imaging of fixed land sites in support of the fires (Section 2.5). Undersea surveillance and custody belong to the undersea force and its own sensors; the fused picture ingests their tracks, but the network does not claim undersea search.

2.2 Requirements

The JSTN must satisfy, at minimum, the following requirements.

Table 2.1 *The requirements the JSTN must meet.*

Requirement	Description
Wide-area detection	Persistent surveillance of the maritime and air approaches across the five sectors (western, Top End, north-east/Coral Sea, south-east, Indian Ocean island arc).
Joint air and maritime tracks	Simultaneous surface and air track production and classification: surface groups, airborne early warning, combat air patrol, strike aircraft.
Targeting-grade custody	Continuous, fire-control-quality custody of designated targets, with sufficient accuracy and currency to commit and update long-range fires.
Mid-course update	An Australian, jam-resistant weapon datalink to update the fires in flight as the target moves.

Requirement	Description
Fusion and C2	Multi-sensor fusion, cryptographic C2, and disciplined find-fix-track-target-engage-assess management across dispersed shooters and multiple axes.
Latency	Sensor-to-shooter timelines short enough to engage a manoeuvring surface group.
Resilience	Graceful degradation under node loss, electronic attack and cyber attack; rapid reconstitution; an organic targeting fallback independent of the space layer.
Cyber and cryptographic integrity	Zero-trust design under assumed continuous, sub-threshold cyber attack; cross-phenomenology integrity checks on the fused picture; post-quantum cryptography and crypto-agility from first fielding [13].
Releasability	A partitioned architecture carrying a national fire-control tier and a releasable partner tier on the same constellation.
Sovereign control	Australian decision and release authority over what is tasked, what is fused and what is fired; Australian design authority and intellectual property over the sovereign core; and resupply from a line Australia can rely on. Space systems do not require domestic manufacture.

2.3 Design concept: the layers

The JSTN fuses complementary layers, each covering the others' weaknesses, into a single joint picture.

Over-the-horizon radar (the wide-area spine). The JORN network provides persistent, very-wide-area north-facing detection of surface and air movement beyond the line of sight: the standing tripwire that cues the finer sensors [4].

The JSTN space sensing layer (SAR shell and ELINT/datalink shell). Australian-owned satellites in proliferated low-earth orbit, providing all-weather wide-area maritime detection and passive emitter geolocation. This is the heart of the network and is developed in full in Section 2.4.

Airborne early warning (the air-picture layer). The E-7A and its successors provide the high-fidelity air picture and battle management, contributing the air half of the joint track and the elevated sensing that defeats low-altitude approaches [1].

The sovereign core (fusion, C2 and the weapon datalink). The software that fuses the layers into a single joint picture, manages the engagement across dispersed shooters, and carries the mid-course updates. This is the part that turns sensing into fires cueing, and the part that must be most rigorously Australian (Section 2.6).

2.4 The space layer: two shells, two procurements

The satellites are a purchase; the software is the sovereign programme. Both shells are bought from mature allied product lines and operated by Australia. A complete national system of this kind has been delivered from contract to operational handover inside twelve months [10], where a first-of-type Australian spaceborne SAR would be a decade-long development programme competing for the same money and workforce as the fires line.

Sovereign control in the sense required in Section 2.2 turns on who tasks the constellation, holds the data and runs the ground segment, not on where the spacecraft was manufactured, and none of those can be switched off by anyone else once the system is owned. What Australia builds is the sovereign core (Section 2.6); the replenishment cycle every five to seven years is the on-ramp for Australian content once there is an operating system to grow it into.

Wide-area sensing and continuous connectivity have different physics, and resolve into two shells at different altitudes and inclinations, procured separately.

The SAR shell

A synthetic-aperture radar (SAR) constellation of the order of sixty small satellites in the ~100–200 kg class at ~520–580 km, bought from an allied supplier and operated by Australia. The current commercial generation is the specification, not an extrapolation from it: the Finnish supplier ICEYE has more than seventy satellites in orbit and has delivered four national systems to European governments inside the last eighteen months [5][10].

The shell does two jobs with the same SAR, and the design states them separately because they are different physics [14]. In wide-swath scan mode (a 200 km swath at ~27 m resolution, ample for ship detection, where surveys at comparable resolution detect better than nine in ten vessels over 50 metres in the open ocean [15]) it searches. At the planning energy budget of 5–10 per cent SAR on-time per orbit, about 70 per cent of it allocated to search, the fleet sweeps the entire surveillance area of the order of every five to ten hours. On-time is not published by any commercial operator and is stated here as an explicit assumption. In strip and spotlight modes (3 m down to 25 cm) it holds custody: any designated point is within a steerable look roughly every 5–16 minutes across the fleet, and a 30-knot warship needs about two hours to cross a single wide-mode frame, so a cued track cannot outrun the geometry.

The sweep is the worst case, not the headline. The ELINT/datalink shell's passive radio-frequency payload watches continuously and cues the SAR onto any emitting contact. That cue travels through the ground rather than through orbit, because purchased SAR satellites carry no optical terminal compatible with the ELINT/datalink shell, so cue-to-look latency is set by the tasking uplink and is tens of minutes rather than minutes; an optical bridge between the shells is a natural step at the first replenishment, and nothing here depends on it. Only the contact that never transmits, a group that has blinded its own air defence to stay dark, waits on the hours-scale sweep. Detection in minutes for emitters and hours for the silent; custody in minutes; identification on demand within a pass.

Orbit design is the lever that is genuinely Australia's to choose, even on purchased spacecraft. Mid-inclination planes (~45–55°) concentrate dwell across Australia's span of latitudes rather than the poles, while the search-weighted planes fly a dawn-dusk sun-synchronous orbit, where near-continuous sunlight buys SAR on-time: a dawn-dusk RADARSAT satellite images 15 minutes per orbit where a satellite of the same mass on a conventional orbit images two [16]. The SAR is the payload that sets the bus, being active and emitting, with the highest power, thermal and antenna demands; but in custody modes it is duty-cycled, so average power fits a small bus, as commercial operators demonstrate across the 85–200 kg class [5]. The cadence figures are illustrative planning estimates from a constellation model (±~30%), not guarantees.

The same shell watches fixed sites, which is developed in Section 2.5.

The ELINT/datalink shell

A larger constellation, of the order of fifty to sixty-six cross-linked satellites at ~780 km, sized for continuous coverage rather than revisit. Continuous coverage is a fundamentally larger job than SAR revisit, which is exactly why it is a separate shell. The constellation model puts continuous coverage of the Australian approaches at roughly forty-four to sixty-six satellites of this class, and the fifty-to-sixty-six planning band used here sits inside that range with margin over its floor.

Iridium is both the existence proof and the likely basis: sixty-six satellites in six planes at 780 km, cross-linked, delivering continuous global narrowband coverage in globally coordinated L-band spectrum [2]. Its acquisition by a vertically integrated launch-and-spacecraft prime, announced June 2026 and expected to close in 2027, places the operating network, the satellite manufacture, the launch vehicle and an optical-terminal business under one roof, with a stated intent to build the next-generation constellation [3]. A derivative of that design, built for and owned by Australia, is the shortest credible path to this shell: a proven architecture, a production line that will be running anyway, and a supplier for whom constellations sold to allied governments are a declared market. It is a candidate path rather than a settled selection, and a competitive acquisition would test it, but it is the one that makes the schedule.

The shell carries three payloads: the passive radio-frequency and electronic-intelligence sensor, the communications relay, and, the load-bearing function, the Australian mid-course targeting datalink. Two of those are the strongest candidates for Australian design authority rather than purchase (the electronic-intelligence payload, and the targeting waveform with its cryptography) being the elements most likely to attract export restriction and the ones whose behaviour must be changeable without asking (Section 5).

Why the payloads split as they do: physics, not preference. The radar cannot share a satellite with a sensitive passive receiver. It is a high-power broadband emitter and would swamp the receiver beside it; bolting a passive listener to a bright emitter defeats the point of both. So the radar anchors its own shell, and the radio-frequency payload lives on the ELINT/datalink shell, where it is passive-compatible, benefits from the satellite count and the precisely known cross-link geometry that emitter geolocation needs (time- and frequency-difference-of-arrival across several satellites), and sits alongside the cueing and datalink it feeds.

Electro-optical and infrared imagery is deliberately omitted from the Australian build and drawn from the commercial market [8]: it is daylight- and cloud-limited augmentation rather than custody, it is the most commoditised and least monopolised space layer and therefore the safest to source commercially.

Two satellite shells, two suppliers, one ground segment

The two satellite shells are bought from mature commercial providers.

The commonality that matters is on the ground: one tasking interface, one sovereign core, one cryptographic regime, and one mission training pipeline even where the platform pipelines differ. The satellite is the commodity. The interface is the asset, and it is Australian (Section 2.6).

The cost of this should be stated: two spares pools, two supplier ground systems to integrate, two sustainment contracts, two configuration baselines. It is a modest and permanent overhead set against a decade of schedule and the first-of-type risk avoided, and both product lines are sustained commercially whether or not Australia buys them.

Frequencies and interoperability

The shell uses three bands by function. L-band (~1–2 GHz) carries the assured targeting and cueing downlink to the weapons: robust in weather, low-power, and closable to the small antenna a munition can carry, for the same reasons Iridium uses it. Ka-band (~20–28 GHz) carries the high-rate imagery downlink to ground. Optical links carry the on-orbit mesh within the ELINT/datalink shell.

Reaching the shooters is two problems, not one, and conflating them overstates what a single bearer can do. The munition is the easy case, because Australia designs it: a purpose-built receiver at the shell's own L-band frequencies, with no legacy constraint to satisfy. A crewed aircraft is the harder case, because its radios are already fitted and certified. A fast jet carries several L-band apertures (its tactical data link terminal and tactical air navigation at 960–1215 MHz, identification at 1030 and 1090 MHz, satellite navigation at 1227 and 1575 MHz), but none of them is a receiver for a mobile-satellite-service waveform 400 MHz higher up, and its satellite communications path is ultra-high frequency rather than L-band. Fitting a new terminal to a fast jet is an integration and certification programme, not a software change.

So leave the cockpit alone and carry a tactical-data-link payload on the ELINT/datalink shell, broadcasting in the 960–1215 MHz band the aircraft's existing terminal already listens on. This has been demonstrated from low earth orbit, the design intent being that terminals already in service receive it unmodified [17]. It also makes the interoperability claim real rather than notional: the shell speaks the grammar the coalition already uses, in the band it already uses, so a JSTN track or cue drops into the existing picture without translation.

2.5 Imaging reconnaissance and land-attack scene matching

The same shell that holds custody of ships is an imaging reconnaissance system over land. The products are catalogue products, published by ICEYE and available today [5]:

Table 2.2 *SAR imaging modes: resolution, scene size and purpose.*

Mode	Resolution	Scene	What it is for
Scan and Scan Wide	15 m and 27 m	100 × 100 km up to 200 × 600 km	Search: finding ships and detecting activity across an area
Strip	3 m	30 × 50 km	Surveying a site and its surrounds; detecting change over large areas
Spot	to 50 cm	to 15 × 15 km	Identifying what is there: buildings, revetments, hardstands, aircraft on an apron
Dwell	25 cm, 50 cm, 1 m	5 × 5 km	Aimpoint-grade detail on a single site

All of it works through cloud and at night, which is the property no optical system has.

Aimpoint development. At 25 to 50 centimetres a fixed site resolves into its parts: which hangar, which revetment, which face of a building, which of several apparently identical structures is the one in use. That level of detail is what lets a missile be aimed at a chosen part of a site rather than at its coordinates.

Change detection. Because the orbits repeat, satellites can be phased so that a site is re-imaged from an identical geometry and successive images compared. SAR is coherent, so the comparison is not of pictures but of phase, and it reveals disturbance far below the resolution of either image: vehicle tracks across bare ground, a revetment newly dug, an aircraft moved on an apron, fresh excavation. This is the mechanism behind three separate jobs (indications and warning, aimpoint currency, and post-strike assessment) and it costs nothing beyond tasking and phasing on the shell already described.

Reference scenes for the seeker. The missile's land attack terminal accuracy comes from a passive machine-vision seeker that matches what it sees against reference scenes held in Australia. Those references are optical, because the seeker is optical, and they are built in peacetime from commercial and national imagery and curated here. The SAR shell's contribution to that chain is not the reference itself but its **currency**: the assurance, obtainable in any weather and on demand, that the scene the seeker will match against is still the scene that is there. Stating the two separately keeps the claim precise, and it identifies the growth path: matching a SAR image directly against an optical reference would let the shell supply terminal-grade references anywhere without a peacetime library. That is an active research field with credible published methods, and the natural scope of a national programme with the Defence Science and Technology Group and the universities: modest in cost, dual-use with civilian earth observation, and publishable. The architecture is positioned to absorb it; nothing here depends on it.

2.6 The sovereign core: fusion, C2 and the weapon datalink

The sovereign core is the one element of the JSTN that must satisfy all three senses of sovereign control at once, and it is the one element no supplier could sell Australia in a form that would still be Australia's.

What it comprises: multi-sensor fusion across over-the-horizon radar, the two space shells, airborne early warning, and shipborne and airborne feeds; track management and custody, including the cross-phenomenology integrity checks that make the picture hard to falsify; engagement management across dispersed shooters on multiple axes; the mid-course update service; the cryptographic regime; and the partition between the national fire-control tier and the releasable partner tier (Section 2.2). It is also the integration point for two suppliers' ground segments (Section 2.4), which is what makes the decision to buy the satellites safe rather than merely cheap.

This is a software programme, not a spacecraft programme, and that is the point. It sits squarely inside what Australian industry can already do, and it is where the enduring value accrues: the constellations will be replaced twice over the network's life, while the core is the asset that persists and appreciates across those generations.

The economics of a programme of this class have also changed. AI-assisted software development means a system of substantial size can be produced, and more importantly continuously re-produced, by a comparatively small team. That matters less for the initial build than for everything after it. It is what lets the software keep pace with the sensors it fuses, absorb a new payload or a new feed as a routine release, and stay on a rolling baseline rather than freezing at contract. It is the same property the resilience posture in Section 2.7 depends on, since a system patched on a rolling baseline presents a moving target where one certified once presents a fixed one, and it is the continuous-iteration idiom of the magazine applied to the network.

AI assistance moves the cost rather than removing it. The binding constraint on software that holds fire-control authority is assurance: verification, test evidence, and the certification that a weapon may be released on this track. That burden does not shrink because the code was written faster, and a programme that assumes it will is the one that fails. What the method buys is throughput and, above all, a sustained rate of change, which is the property the architecture actually needs and the one conventional acquisition has been least able to deliver.

2.7 Resilience and reconstitution

Because the JSTN is the centre of gravity, an adversary's sharpest move is to blind it. The design answers this by being hard to blind and by making the attempt self-defeating.

Hard to blind. The fixed nodes are dispersed, hardened, redundant and designed for rapid reconstitution, so node loss degrades fidelity rather than severing function, and the multi-layer design means no single sensor type is a single point of failure. The bearers degrade rather than fail: if the high-rate relay is jammed or lost, an assured narrowband layer still carries the tracks, cues and fire orders, which are hundreds of bytes rather than megabytes, so the network loses data rate rather than going dark. The survivable object is the constellation, not the satellite; resilience comes from numbers and rapid replacement, not from gold-plating each spacecraft.

Self-defeating to attack. The JSTN draws on Australian, allied and commercial multinational feeds, so striking its Australian nodes is open war on the Australian homeland, and striking its commercial feeds internationalises the conflict. Every physical route to blinding it crosses an escalation threshold, which itself contributes to deterrence. Commercial feeds and bearers are a resilience complement, not the assured layer; the Australian nodes remain the backbone.

The cyber front: the one blinding route that is not escalatory. The escalation logic above has an exception, and the design must own it. Cyber attack is deniable, sub-threshold and already continuous; the coercion campaign at Section 1 opens with it. The network is therefore designed for standing attack, not peacetime sanctuary: a zero-trust interior in which every node, feed and operator is continuously authenticated; segmented, cross-checked feeds, so a poisoned input must survive SAR, radio-frequency and over-the-horizon-radar analysis simultaneously, since the fusion that makes the picture accurate is also what makes it hard to falsify; and the assumption that any fielded missile can be captured, so no airframe carries a secret that unlocks the fleet. AI weaponisation sharpens this threat rather than changing its nature: machine-speed vulnerability discovery will probe the attack surface faster than any static defence can be certified. The credible posture is the one the sovereign core is built for (Section 2.6), which is continuous iteration across the fleet, with the same machine-speed methods turned to defensive monitoring.

2.8 Building on existing investment

The JSTN is achievable because it extends and integrates capability the ADF largely already fields. The decisive new investment is the integrating layer and the space-custody layer, not the sensors, most of which exist.

Table 2.3 *Existing Australian investments and what each contributes to the JSTN.*

Existing investment	Contribution to the JSTN
Jindalee Operational Radar Network	Wide-area joint air and maritime detection spine
E-7A airborne early warning and control	High-fidelity air picture and battle management
Integrated air and missile defence C2	The joint air-defence C2 to fuse with
Defence and commercial ISR and satellites	Space-based sensing, augmented by an Australian maritime-custody layer
New: the sovereign core	**Fusion, C2 and the weapon datalink: turns the joint picture into fires cueing**

The institutional point. The JSTN should be developed as joint-force infrastructure with the fires hung off it, not as a bolt-on. The 2026 Integrated Investment Program strengthens the footing, funding 'the targeting enterprise' alongside long-range strike as a named priority of the strategy of denial [12]. The ask is not a new category of investment but the completion, at decisive scale, of one the Government has already opened.

2.9 Sizing and cost: the trade space

What the space layer costs is not a single number but a small trade space across both shells, their satellite counts, and the coverage each buys. The anchor is not an estimate but the market.

The cleanest published price for a complete national system is Poland's POLSARIS: about €200 million for four satellites plus fixed and mobile ground segment and operator training, contracted May 2025 and handed over operational within about twelve months, with Finland's ~€158 million three-satellite programme corroborating the scale and the Netherlands, Portugal and Greece buying similar fleets at undisclosed prices [10]. Germany's ~€1.7 billion five-year contract with the same supplier is the contrast that prices the difference between owning a system and renting a service [10].

Two lessons from those transactions price the shell. First, most of the money is not the spacecraft: the satellite itself manufactures for single-digit millions while the delivered system prices some twenty times that, so the cost model is a fixed sovereign segment (ground stations, tasking, the sovereign core, training and operations stand-up, of the order of A$0.3–0.5 billion once) plus a marginal cost per satellite delivered and launched at production volume, of the order of A$25–40 million. That band is derived from the Polish and Finnish transactions as an explicit planning figure, not a quoted price. Second, small-fleet system prices cannot be multiplied linearly to sixty satellites, because the fixed segment amortises and the production line is precisely the procurement-model logic (Glossary, GOCO).

Against those anchors, the design space is:

Table 2.4 *The space layer sized and costed: the trade space (A$, illustrative).*

Shell / option	Coverage	Cost (illustrative)	Character
Sensing, ~60 satellites (baseline)	National; custody on cue ~5–16 min (access); full-area detection sweep ~5–10 h	~A$2.0–3.0 bn	All-weather search, custody and identification of the approaches and the whole continent
Targeting / communications, ~50–66 satellites	Continuous	~A$1.2–1.8 bn	The Australian fire-control link: cheaper satellites, more of them, optical cross-links within the shell
Combined space layer	—	~A$3.2–3.8 bn	

For scale, the whole space layer costs less than a single cancelled Australian programme: JP 9102, the geostationary SATCOM system estimated at up to A$7 billion before Defence cancelled it in November 2024 in favour of exactly the multi-orbit, proliferated posture designed here (the policy ground is quoted at Section 5.2) [11].

2.10 Development sequence, and the interim

Consistent with the architecture's progressive build, the JSTN is fielded incrementally, and each increment adds usable targeting capability. Partial completion is partial capability, so the network deters earlier and improves continuously rather than waiting on a single delivery.

1. Fuse what exists, and start the core. Integrate the over-the-horizon radar, airborne early warning and IAMD C2 into a fused joint picture, and stand up the sovereign core against that picture and against commercial feeds already under contract. The core starts first because it takes longest to mature and because it is the thing that must be Australia's.

2. Buy an initial SAR shell. On the delivered benchmark, a first national SAR capability is a near-term deliverable measured in a small number of years rather than a decade [10]. It arrives operational, with its ground segment and operator training, and begins contributing custody immediately.

3. Grow the SAR shell and procure the ELINT/datalink shell. Scale the SAR constellation toward its full size while the ELINT/datalink shell, the electronic-intelligence payload and the Australian targeting datalink mature behind it. The sovereign core is the long pole and is sequenced accordingly.

4. Replenish, and grow the Australian content. The first replenishment cycle is the natural on-ramp for Australian content, and for the optical bridge between the shells (Section 2.4).

Subscribe to commercial SAR data from the outset. A tasking subscription with the constellation supplier is available immediately and at a small fraction of the shell's cost, and it does two things. It provides SAR custody over the approaches years before the owned satellites fly. And it gives the sovereign core a live operational feed to be built and tested against, in the products and formats the owned shell will later deliver, so the core matures on real data rather than on simulation and is already integrated at handover: buying from the eventual constellation supplier makes the interface the core is written to the interface it keeps. The subscription is a service and not a capability, and tasking priority remains the supplier's in a crisis, so it is time-boxed against the owned shell and retires into the backup layer as that shell comes up.

The interim, and the permanent backup. Until the owned shells are up, and afterwards as a resilience layer, allied and commercial operators already provide every layer today: SAR [6], radio-frequency geolocation [7], electro-optical imagery [8] and satellite communications [9]. The governing principle is diversification: no single foreign provider should become a dependency, and the owned layer displaces the rented one as it matures.

Table 2.5 *Sensing and communications layers: representative allied and commercial providers, and the acquisition path for each.*

Layer	Representative allied and commercial providers (illustrative)	Path
SAR (all-weather custody)	ICEYE (Finland); Capella, Umbra (US); Synspective (Japan)	Purchased and Australian-operated: the SAR shell
RF / ELINT (emitter geolocation)	HawkEye 360 (US); Unseenlabs (France, maritime)	Purchased shell, with Australian payload content
EO / IR (visual identification)	Airbus Pléiades Neo (Europe); Maxar/Vantor, Planet, BlackSky (US)	Commercial: deliberately not built; daylight- and cloud-limited augmentation
Communications	Starlink / Starshield (US); OneWeb–Eutelsat (Europe); Viasat–Inmarsat (L-band); Iridium (narrowband)	Australian relay and targeting datalink via the ELINT/datalink shell; commercial as diversified backup bearer

The providers are named to establish that each layer is real and procurable today from allied sources, not to endorse any one. Capabilities, ownership and availability change, and several are subject to export control or supplier-government priority in a crisis, which is precisely why they are the interim-and-backup layer and the owned shells are the assured layer.

References

All references were accessed in July–August 2026. Several concern fast-moving programmes and should be re-verified before external use. Technical values are drawn from public specifications and are cited to establish feasibility, not to fix a design.

[1] Boeing Defence Australia / RAAF, E-7A Wedgetail airborne early warning and control. An in-service, Australian-operated AEW&C aircraft, cited as an existing sensing and battle-management node of the JSTN.

[2] Iridium Communications technical references. The constellation operates in low-earth orbit (~780 km) with an L-band service (1616–1626.5 MHz); Iridium Certus is narrowband (~704 kbit/s down / ~352 kbit/s up) but globally covered, low-latency and comparatively hard to deny: an assured cueing-and-command bearer rather than a sensor pipe.

[3] Rocket Lab Corporation, 'Rocket Lab to Acquire Iridium', 29 June 2026 (~US$8bn enterprise value; close expected 2027; stated intent to support development and deployment of Iridium's next-generation constellation, and access to globally coordinated L-band spectrum); 'Rocket Lab Completes Mynaric Acquisition', 14 April 2026 (laser optical communications). Cited for the vertically integrated launch, spacecraft, optical-communications and network base from which an Australian-owned derivative of the ELINT/datalink shell could be procured.

[4] Australian Department of Defence / BAE Systems Australia, Jindalee Operational Radar Network. Australian over-the-horizon radar providing wide-area surveillance of the northern and western approaches: the existing sensing spine the JSTN fuses.

[5] Small-satellite SAR design and imaging references: ICEYE X-band constellation (Finland), 85 kg (Gen2) to ~200 kg (Gen4); published imaging modes per ICEYE product documentation: Dwell at 25 cm, 50 cm and 1 m over 5 × 5 km; Spot to 50 cm over up to 15 × 15 km; Strip at 3 m over 30 × 50 km; Scan and Scan Wide at 15 m and 27 m over 100 × 100 km up to 200 × 600 km; Capella Space (US), ~100–190 kg. eoPortal; ICEYE product documentation and press releases, 2025–26. Cited for the mode set in Section 2.5 and to establish that capable wide-swath and spotlight SAR is achievable from the ~100–200 kg class. No commercial operator publishes SAR duty cycle or imaging time per orbit; the on-time figures in Section 2.4 are explicit planning assumptions, not citations.

[6] Commercial SAR operators: ICEYE (Finland; more than 76 satellites launched by mid-2026); Capella Space and Umbra (US); Synspective (Japan, StriX). Trade and market reporting, 2026. Cited as the interim and backup all-weather detection layer available from allied sources.

[7] Commercial RF-geolocation and ELINT operators: HawkEye 360 (US; clusters of three formation-flying satellites, 14 clusters and 42 satellites deployed by early 2026; software-defined receivers tuning 70 MHz to 6 GHz, extended to about 18 GHz with downconverters; contracted to the US National Reconnaissance Office under its Strategic Commercial Enhancements arrangement, extended December 2025; supplied to Australia's Department of Foreign Affairs and Trade from July 2023, delivering maritime domain awareness to the Pacific Islands Forum Fisheries Agency under the Quad's Indo-Pacific Partnership for Maritime Domain Awareness). Unseenlabs (France; BRO maritime radio-frequency constellation, 21 satellites launched by May 2026; single-satellite direction-of-arrival geolocation at approximately 5 km mean accuracy; shortlisted for France's national CERES successor). Company, government and trade reporting, 2023–26.

[8] Commercial electro-optical operators: Airbus Pléiades Neo (Europe, 30 cm, intraday revisit); Maxar/Vantor WorldView Legion (US, 30 cm-class); Planet (US, SkySat/Pelican); BlackSky (US, Gen-3, mid-inclination orbit for time-diverse revisit). Trade reporting, 2026.

[9] Government and military satellite communications: SpaceX Starlink and the Starshield variant (encrypted, jam-resistant, optical inter-satellite links); OneWeb–Eutelsat (Europe, LEO, naval trials); Viasat–Inmarsat (GEO L-/Ka-band, maritime, merged 2023); Iridium (LEO L-band narrowband). SpaceX, SpaceNews and Payload, 2025–26. Cited to show diverse allied bearers exist; the sovereignty caveat in the text applies.

[10] National SAR procurement benchmarks: ICEYE POLSARIS (Poland), ~€200m for four satellites plus fixed and mobile ground segment and operator training, contracted May 2025, three satellites operational within ten months and the system handed over to the Polish Geospatial Reconnaissance and Satellite Services Agency in May 2026, with Polish operators trained to run the constellation independently; Finnish Defence Forces, ~€158m for three satellites plus options, September 2025; Royal Netherlands Air Force and Portuguese Air Force fleets and Greece's programme (values undisclosed); Rheinmetall–ICEYE contract with Germany, ~€1.7bn over five years, a data-and-operations service rather than a system purchase; Swedish Armed Forces, ~SEK 1.3bn across ICEYE (SAR) and Planet (optical) combined, vendor split undisclosed. ICEYE, Rheinmetall and government press releases, 2025–26. Cited for the delivered-system price, the ownership-and-operation model, and the timeline; the ~A$25–40m marginal per-satellite band and the fixed-segment estimate in Section 2.9 are derived planning figures, not quoted prices.

[11] JP 9102 (Australian Defence SATCOM System): Joint Project 9102 Ph.1, awarded to Lockheed Martin Australia April 2023 for 3–5 geostationary satellites plus ground systems, estimated ~A$5.9–7bn; cancelled 4 November 2024, Defence citing the vulnerability of a single-orbit approach and a shift to a resilient multi-orbit capability. Australian Defence Magazine, SpaceNews, Breaking Defense and The Strategist (ASPI), November 2024–2025.

[12] Australian Government, Department of Defence, '2026 National Defence Strategy' and '2026 Integrated Investment Program', released 16 April 2026. The IIP names the targeting enterprise, long-range strike, IAMD and uncrewed systems as funded priorities of the strategy of denial.

[13] NIST, FIPS 203, 204 and 205 (ML-KEM, ML-DSA and SLH-DSA post-quantum cryptographic standards), finalised August 2024; Australian Signals Directorate guidance on planning the transition to post-quantum cryptography. The basis for the first-fielding post-quantum requirement in Section 2.7.

[14] A. Bojarski and M. Bachmann, 'A Simplified Approach to Estimate the Area of Interest Coverage Duration Performance for SAR Satellites', IEEE Transactions on Aerospace and Electronic Systems, 2024. The formal distinction between access revisit and area-coverage duration relied on in Section 2.4.

[15] F. S. Paolo, D. Kroodsma, J. Raynor et al., 'Satellite mapping reveals extensive industrial activity at sea', Nature 625, 85–91 (2024). Satellite SAR at ~20 m resolution detected more than 90 per cent of vessels over 50 m across 753,030 images. The basis for the detection-rate figure in Section 2.4.

[16] Canadian Space Agency, RADARSAT Constellation Mission technical characteristics: 1,400 kg satellites on a dawn-dusk sun-synchronous orbit, 15 minutes of imaging per orbit per satellite, orbit-average power under 220 W. The dawn-dusk duty-cycle existence proof. Contrast: KOMPSAT-5, ~1,400 kg on a conventional sun-synchronous orbit, 2 minutes per orbit; eoPortal.

[17] Tactical data link from low earth orbit: US Space Development Agency, first Link 16 network transmission from space to ground, demonstrated 21–27 November 2023 on Tranche 0 Transport Layer satellites, with subsequent demonstrations to airborne and shipborne terminals; the programme's stated design intent is reception by terminals already in service. Regulatory constraint: US Federal Aviation Administration restrictions prevented broadcasting into the national airspace system, and testing was confined to international

waters and allied territory under a national telecommunications waiver. Space Development Agency releases 2023–26; The War Zone; Air & Space Forces Magazine. Cited to establish that the path exists and that frequency clearance, not the radio, is the pacing item.

Section 3 — Long-Range Precision Fires: Requirements and Design Concept

If the JSTN is the half of the architecture that finds, the Long-Range Precision Fires are the half that kill. This section states what the fires must do, then specifies the missile that does it, how it is launched and how it flies. The effectiveness model is at Section 4; the industrial path is at Section 5.

3.1 Requirements

Five requirements set the design. They are stated first because the missile follows from them, and because between them they rule out buying one.

Table 3.1 *The five requirements, what each demands, and what each drives in the design.*

Requirement	What it demands	What it drives
Very long range	Hold a coercing force at risk from dispersed Australian sites, from beyond the reach of the missiles it would trade against, and extend deterrence forward of the continent. Chinese land-attack cruise missiles of the CJ-10 class are publicly assessed at 1,500–2,000 km [6].	A design point of **2,400 km**, which keeps its margin even if those public estimates prove low.
Affordable mass	Salvos large enough to exhaust the interceptor magazine of a highly defended group, which makes unit cost a critical requirement.	The **US$300,000 class**, produced in volume: 144 rounds to a battery, 3,600 to the force.
Lethality against maritime and land targets	A single hit capable of destroying a major surface combatant, and the same round useful against land targets, air defences and materiel.	A **250 kg** blast-fragmentation warhead, three fuze settings, and a seeker that chooses the aimpoint.
Deployability, low maintenance and support	Deploy to remote and forward locations on existing roads and sit there for long periods with minimal support.	Launch from a **standard 40-foot high-cube ISO container**; no hydraulics, no fixed infrastructure, small battery crew.
Centralised control	Fire on Australian authority as part of collaborative, networked and joint missions, including simultaneous time on target with other shooters.	A **command datalink with the JSTN** for cueing and mid-course update, and time-of-arrival synchronisation.

Nothing in or entering service meets the first two together. The longest-reaching off-the-shelf land-launched option reaches about 1,000 km [2]; the missile Australia is already building for maritime strike reaches about 500 km [1]; and both sit near US$2 million apiece, which is a point-strike price, not a saturation one. Section 5 sets the two production models against each other and prices the gap.

The last two requirements decide who builds it. Resupply that cannot be gated by another government's export approvals or wartime demand, and a roadmap (JSTN integration, software, seeker, every capability increment) set on Australian timelines, are not purchasable properties at any price. They are why this is an Australian design on an Australian line rather than a procurement. That does not mean acting apart from allies: the fires take cueing from allied as well as Australian sensing, and a line at this scale can supply allied magazines rather than compete for them. The industrial path is Section 5.

3.2 Long Range Strike Missile

The fires are built around a single round: the Long-Range Strike Missile (LRSM), a subsonic, low-observable, container-launched cruise missile with a machine-vision terminal seeker.

Table 3.2 *The Long-Range Strike Missile, summarised.*

	Long-Range Strike Missile
Reach	2,400 km assured, container-launched from dispersed Australian sites
Warhead	250 kg blast-fragmentation; point, delayed or airburst fuzing
Guidance	Satellite and inertial navigation, command uplink to the JSTN, passive machine-vision terminal seeker
Profile	High subsonic cruise-climb, horizon-hugging descent, terminal sea-skim
Unit cost	US$300,000 class, produced in volume on an Australian line
Launch unit	One 40-foot high-cube container, nine rounds
Battery	Sixteen containers, 144 rounds, some US$43 million

Mass fires. Each round in a salvo is armed and each is capable of destroying a ship, so a defended group cannot triage the raid: every interceptor it spends is spent on a genuine threat, and there is no track it can safely leave alone. The salvo is its own saturation, which is why the architecture fields no decoy rounds: a decoy buys the defender's attention with something that cannot hurt it, and at US$300,000 a real round buys the same attention and can. The effect is modelled at Section 4.

The same round serves land targets. Precision substitutes for volume of fire against any target whose critical points can be seen and identified, so the LRSM covers land strike and anti-materiel, and, with an anti-radiation seeker in place of the optical one, the suppression of air defences. Fixed sites are found, aimpointed and assessed from orbit (Section 2.5). A mobile system that neither emits nor stops remains a genuinely difficult problem and is not claimed on this basis.

The design, sized

The design point is derived rather than asserted, from a Breguet range calculation calibrated on a small turbojet, a fixed length budget for nose, warhead bay, engine and aft section with an in-line booster, and a 0.58 m body.

Table 3.3 *The design sized: quantities, values and the basis for each.*

Quantity	Value	Basis
Warhead	250 kg blast-fragmentation	Harpoon class [5]: the weight assessed to destroy most frigates and destroyers
Body diameter	0.58 m	Sized to pack three across in a 40-foot container
Overall length with booster	6.75 m	Fixed sections 3.51 m plus 3.24 m of tankage
Boost section	0.70 m	A solid motor of the Tomahawk's class, 6,000 lbf for twelve seconds, scaled for launch mass and body cross-section. The Tomahawk carries 0.69 m of boost section at 0.53 m diameter, 5.56 m unboosted against 6.25 m boosted [3]
Fuel	804 kg	Fuel fraction 0.49
All-up mass	~1,650 kg	Flying mass at the start of cruise. Launch mass is about 1,885 kg with the boost section attached
Design range	2,400 km	Outreaches the CJ-10 class at ~1,500–2,000 km [6] with margin
Length available in the container	7.43 m	0.68 m still in hand

The design holds about twenty per cent of range in reserve. Flown at its best altitude and speed, in still air and on a direct track, the round reaches about 2,900 km against the 2,400 km it is specified to hold. The specification is the assured figure rather than the maximum: 2,400 km is what the round delivers against a 30 m/s headwind at cruise altitude on a track routed around neutral airspace. The difference between the two is reserve, and it is spent on wind, routing and fuel held for the approach.

Range is a policy decision. Tankage alone, taken to the full length the container allows, reaches about 3,100 km; a small turbofan in the same container, launcher and battery would be of order 4,000 km. Where the reach stops determines which targets the force can be read as holding at risk, so it is set by Government against declared deterrence policy, not by what the airframe will carry.

Engine and airframe are chosen for cost. A turbojet rather than a turbofan: fewer parts and stages, cheaper at scale, at the price of a larger fuel fraction, which is acceptable for a missile produced in thousands. The small jet engine was once the hard sovereign component and is not any longer. It is built in-house by commercial makers in this class today [4], and additive manufacture is driving cost and lead time down further. The airframe is whatever is cheapest at volume for a low-observable shape with adequate structural margin, not the capital-intensive process used on crewed aircraft.

Warhead: 250 kg of blast-fragmentation, with three fuze settings. *Point* detonates on impact, where the aimpoint is the structure struck. *Delayed* passes through outer plating before detonating, where the aimpoint is a space rather than a surface: a command centre, a machinery space, a magazine. *Airburst* detonates above the target and is the setting for soft, dispersed and unarmoured targets (launcher revetments, parked aircraft, vehicle parks, fuel and logistics) where fragmentation across an area does more than a crater. The fuze is the only difference between a suppression round, a ship-killer and an area-attack round: one charge, one warhead line, one stock.

Aimpoints: the ship's critical spaces. The seeker is aimed at what ends the ship's fight. First the vertical launch magazines, where a delayed-fuze detonation among loaded cells starts a propellant cascade the ship cannot fight: the mechanism by which one hit becomes a ship-kill rather than a mission-kill, and the widely assessed cause of the loss of HMS Sheffield in 1982 [7]. Then the command spaces, then machinery and propulsion. Sensors, directors and launchers present first on a closing approach and will take some hits regardless; the model counts those as losses rather than as a phase, which is the conservative reading.

Navigation, command and the terminal seeker. Satellite and inertial navigation, with a command uplink to the JSTN for in-flight retargeting as the group moves. Terminal accuracy comes from a passive optical and infrared machine-vision seeker that recognises the target and selects the aimpoint. Because it is passive and optical it cannot be jammed the way a radar seeker can, and because it is independent of the satellite links, accuracy survives the loss of both navigation and network. That seeker is the design's load-bearing element and the component with the steepest sovereign improvement curve (Section 5).

3.3 Launch: the 40-foot high-cube container

The launch unit is a standard 40-foot high-cube ISO container on a side-loading trailer. The arithmetic is derived from published ISO dimensions [8] rather than assumed, because launch density sets the battery and the battery sets the force.

A 40-foot high-cube is 12.032 m long, 2.352 m wide and 2.693 m high internally. Allowing for the support frame, the deck and the roof mechanism leaves 2.192 m of usable width and 2.443 m of usable height. Each round sits in a 0.68 m tube.

Table 3.4 *Container fit: required against available, with margins.*

Check	Required	Available	Margin
Three tubes across the width	2.04 m	2.192 m	0.152 m
Three tubes up the height	2.04 m	2.443 m	0.403 m
Round length against length available to the round	6.75 m	7.43 m	0.68 m

Three across by three high is **nine rounds per container**, in a 3 × 3 eggbox. The high-cube is required rather than preferred: a standard-height container leaves about a hundred millimetres for frame, rack and tolerance across the whole stack, and the extra 0.30 m converts a marginal fit into a comfortable one at no cost in handling. Elevation is by electric motors on a gear rack rather than a central hydraulic ram, which frees the centre cell (the difference between eight rounds and nine) and removes a hydraulic power pack from a container that may sit unattended for months.

Elevation and efflux. The roof opens and the stack elevates to roughly 35 degrees: enough for a boosted launch, low enough to keep the centre of gravity inside the container footprint without outriggers. Efflux vents up through the open roof and aft through a blow-out panel in the back wall.

Length and mass. Of the 12.032 m internal length, 7.43 m is available to the round against 6.75 m used. The 6.75 m includes the 0.70 m boost section, which is the length the container has to hold; the airframe that flies is 6.05 m. Nine rounds at about 1,885 kg at launch is 17.0 tonnes of ordnance; with tubes, frame, gear and the container itself, gross mass is of order 25 tonnes against an ISO maximum of 30,480 kg. Length is the binding constraint, not mass, and it has margin in it. A loaded container is at the upper end of standard container handling, so Army will need prime movers and side-loading trailers matched to the 40-foot high-cube at that gross mass.

Power and environmental control. The container is the magazine as well as the launcher and rounds may sit in it for years, so internal temperature and humidity are actively controlled: propellant, warhead, batteries and seeker electronics age at the rate their storage conditions set. The roof carries a solar array on the leaves that open for launch, and insulation with deployable side shades reduces the heat load the plant must remove; regulation is by reverse-cycle air-conditioning, which dehumidifies as well as cools. Lithium iron phosphate batteries buffer the array and carry the duty cycle for communications, electronics and conditioning, the chemistry chosen for thermal stability and cycle life beside ordnance. Monitoring and anti-tamper measures erase symmetric key material and sensitive software on interference. A conditioned container has a thermal signature, so decoy containers carry the same fit-out.

Deployment is static and dispersed. There is no shoot-and-scoot concept, and that is a decision rather than an omission. Containers deploy dispersed across a sector on Defence land, with empty decoy containers alongside them as targets; basing, manning and site preparation are in Section 4.4. Destroying a battery is therefore a substantial investment in long-range fires, or in platforms that must be brought well forward to deliver them.

Forward basing extends the reach. From forward sites the area held at risk grows sharply; Manus Island in particular sits forward in the Bismarck Sea and would project coverage across approaches no mainland site can reach. Dispersed pre-positioned fires are more defensible forward than exquisite assets, because losing a container of rounds to a pre-emptive strike is tolerable where losing a ship or an airbase is not. Manus is sovereign PNG territory, so the option is contingent on host-nation agreement.

3.4 The attack profile

The raid is a network engagement with a kinetic conclusion. Before any interceptor flies, the contest is between the raid's custody-and-replanning loop and the group's warning-and-picket loop. The attack profile is the physical expression of that: the missile flies the geometry the network hands it.

The salvo arrives together. Missiles launched from dispersed containers are timed to converge on the group simultaneously and from divergent bearings. A layered defence performs as designed when it can fight in sequence, servicing one set of tracks and recovering its fire-control channels before the next arrives; a salvo that arrives at once loads magazine, fire-control capacity and crew together. Time-of-arrival synchronisation across dispersed launch points is therefore a design requirement on the navigation suite, not a tactical nicety.

A horizon-hugging descent, not a cruise-then-pop. The missile cruises at about Mach 0.85, climbing slowly from about 8,000 m to about 12,000 m as fuel burns off, which holds it at the altitude and speed where a turbojet's range per kilogram of fuel is greatest. It then begins a gradual descent from a couple of hundred kilometres out that tracks the radar horizon of the outermost ship in the group's disposition downward: high and fast when far out, progressively lower as it closes. Against a masthead radar the permissible altitude falls smoothly along the approach, on the order of 1,800 m at 200 km, 350 m at 100 km and wave-top inside 40 km, so the altitude given up late is converted into terminal speed rather than paid for in fuel.

Airborne radar sets a different horizon, and the profile answers it. The descent schedule above is drawn against a masthead. An airborne early-warning radar at 9,000 m holds line of sight to a round at cruise altitude out to some 760 to 840 km, and still reaches 410 km against one on the deck, so descending early cuts the geometry roughly in half rather than defeating it. Where a group is assessed to have an airborne radar up, the profile is flown lower and the descent begins further out, at a cost in fuel the range margin covers. Detection inside those bounds depends on the round's signature against a look-down radar; the geometry sets only the ceiling.

Warning is not interception, and the air wing is counted as a magazine. A combat air patrol acting on early warning carries a fixed number of air-to-air rounds. It gets one pass at a salvo timed to arrive together, it cannot land, rearm and relaunch inside that window, and it cannot leave tracks unengaged, because every round it does not kill continues to the group. The limit is the same countable one the surface magazine imposes, and Section 4 counts the air wing's rounds on that basis, at an allowance set deliberately in the defender's favour.

Terminal sea-skim. The terminal is the profile flown by every fielded sea-skimming anti-ship missile. A skimmer stays below the ship's radar horizon until roughly 40 km or less, which guts the long-range area-defence layer, compresses every engagement window and degrades interceptor homing in sea clutter and multipath. A radar altimeter holds a small dynamic margin above the live wave tops rather than a fixed altitude, so the skim is as low as the sea state safely allows: of order five metres in a calm sea, and higher as the sea gets up. The one departure from the fielded pattern is the seeker, which is passive, so nothing radiates on the run in.

The aimpoints are reached from the side. A delayed fuze carries the warhead inboard through the ship's side, and the spaces Section 3.2 lists present there: a vertical launch magazine is a stack of cells seven to nine metres deep below the deck it opens onto, and command spaces, machinery and fuel sit lower again. This is the profile that disabled HMS Sheffield [7]. A carrier is attacked the same way, and its flight operations end when it cannot make wind over the deck, power its catapults, or fuel and arm its aircraft. The land-attack profile is set by terrain and scene reference, and is not developed here.

References

All references were accessed during preparation of this volume in July 2026. Several of the systems cited are in rapid development; a reader revisiting later should re-verify, as figures move. The systems below are cited to establish that the LRSM's component choices are fielded or in production, not that any of them is the proposed design.

[1] Standard references on the Precision Strike Missile (PrSM), the 500 km-class ground-launched missile fired from HIMARS and M270 (roughly US$2M per round). Australia is a cooperative programme participant and has selected PrSM with HIMARS for maritime strike, with a June 2025 co-production memorandum and onshore manufacture planned.

[2] Naval News, 'Japan Deploys New Longer-Range Missiles, Formally Designates Type 25 Systems', 31 March 2026; USNI News, 1 April 2026. Type 25 surface-to-ship missile, ground-launched, about 1,000 km, operational 2026. Cited as the longest-reaching land-launched option in service.

[3] Standard references on the Tomahawk (BGM-109) family: US Navy fact file; Raytheon and RTX product literature; open cutaway and specification data. Cited for the airframe arithmetic at this range and warhead class: the BGM-109A and BGM-109C are one airframe at identical external dimensions, carrying 130 kg to roughly 2,500 km or 450 kg to roughly 1,250 km.

[4] Destinus, Ruta product pages (Block 2, Block 3), destinus.com; with EDR Magazine, 'Eurosatory 2026: Destinus unveiled a family of high-tech strike cruise missiles', June 2026. A container-launched, turbojet-powered, AI-guided cruise missile family in current production, cited to establish that the type and its engine are built commercially today, not as an airframe to licence or acquire.

[5] Standard references on the Harpoon anti-ship missile (WDU-18/B, roughly 221 kg blast-fragmentation warhead), with open sink-exercise reporting on the difficulty of sinking large hulls. Cited for the warhead class assessed to destroy most frigates and destroyers.

[6] CSIS Missile Threat, 'CJ-10 (DH-10)'. Cited for the adversary's ground- and sea-launched land-attack cruise-missile reach (roughly 1,500 to 2,000 km class; public range estimates, not measurements).

[7] Standard historical references on the 1982 loss of HMS Sheffield, widely assessed as a propellant and fire cascade initiated by a strike whose own warhead did not detonate. Cited as the mechanism by which a magazine hit becomes a ship-kill.

[8] ISO 668, series 1 freight containers, 40-foot high-cube (1AAA): 12.032 m internal length, 2.352 m internal width, 2.693 m internal height, 30,480 kg maximum gross mass. Cited for the launch-density arithmetic in Section 3.3.

Section 4 — Fires Effectiveness and Force Structure

This section models a massed, single-salvo, multi-axis raid launched from dispersed containers against three named threats, and derives from that model the rounds required, the battery, the force and the cost-exchange. The missile itself, its launch container and its attack profile are specified in Section 3. The anti-ship case is modelled in detail because it is the most demanding of the missile.

4.1 The effectiveness model

The model counts cells. A defended surface group's air defence rests on a countable quantity: the number of vertical launch cells it carries and the number of point-defence rounds in its launchers. Four facts follow from that, and the model is built on them alone.

1. Every cell is either loaded with an interceptor or it is not. 2. An interceptor fired is spent whether or not it kills. 3. A vertical launch system cannot be reloaded at sea. No navy has a fielded, reliable capability to do it in any but the calmest conditions, and none plans on it. 4. Because there is one round in the salvo and it is real, the defender must engage every track it detects.

From these, and without agreeing on a single kill probability, the raid size that empties the magazine can be counted. That is the point of building the model this way: a cell count is checkable from public sources, where a chain of per-layer kill probabilities is not. A defender can raise its kill probabilities and make the magazine last a little longer, but it cannot make the magazine deeper.

The arithmetic. For each threat case:

1. Maximum kills the defence can achieve is the sum of what every layer could do if every shot were taken and every engagement went as designed: area cells times the area kill probability, plus point-defence cells times the point-defence kill probability, plus a fixed effective-kill allowance for close-in guns and hard-kill directed energy per warship, plus a combat air patrol allowance where a carrier is present.
2. Rounds to exhaust the defence is that number divided by flight reliability. The figure counts every layer, the magazine-limited and the rate-limited alike. At this point the defence has fired everything and the remaining rounds arrive against guns alone.
3. Rounds to destroy adds the hits actually required on the ships, divided by terminal hit probability, and divides the total by flight reliability.

The model is deliberately generous to the defender. It credits every cell as available and loaded with an interceptor. It ignores the fire-control channel limit. It ignores the simultaneity ceiling a single coordinated salvo imposes, which is the mechanism most likely to make the real collapse sharper than the arithmetic. It gives no credit for the degradation caused by the hits themselves, even though a ship that has lost a director fights worse than one that has not. The real requirement is lower than the table says. The purpose is a raid size that does not depend on winning any of those arguments.

The group model assumes no fire-control channel limit. The group is credited with its area cells where Section 4.2 denies them to a lone ship, because dispersed escorts break the horizon at different times and cue one another, so the group's area layer engages where a single ship's cannot; what the horizon takes from the group is reach rather than inventory, since a long-range interceptor fights a thirty-kilometre engagement and every shot is a late shot. With active radar homing interceptors and no separate illuminators on the Type 052D and 055, the defender is bounded by cells and decision time rather than by channels, and the model credits it the whole magazine. This is the defender-favourable reading, and it is deliberate: the group is priced on magazine exhaustion. A single ship in the brief window of a saturation attack is the opposite case (it cannot cycle its channels through its magazine before it is overwhelmed), so Section 4.2 prices single ships on that channel-limited basis. Held to the same basis, the group's raid sizes would fall, not rise.

4.1.1 Parameters

Table 4.1 *Effectiveness model parameters and the basis for each.*

Parameter	Value	Basis
Type 055 area SAM cells	75	≈⅔ of 112 VLS [3]; the group's air-defence command ship. The SAM fraction is the model's principal assumption and is swept in Section 4.1.3. Cell counts are credited as interceptor counts; a deeper loadout is bounded by the fire-control ceiling rather than by the magazine.
Type 052D area SAM cells	43	≈⅔ of 64 VLS; an area-air-warfare destroyer, freed to weight air defence by the frigates' anti-submarine screen
Type 054B area SAM cells	19	≈60% of 32 VLS; balance anti-submarine and anti-ship, since the frigate screens against submarines
Carrier area SAM cells	0	point defence only; the air wing is modelled separately
Oiler (Type 903A)	0	four single-barrel 30 mm mounts, no effective defence against a sea-skimmer
Point-defence cells per warship	24	HHQ-10 launcher
Area SAM kill probability	0.50	sea-skim degraded, at the terminal skim height specified in Section 3.4
Point-defence kill probability	0.47	sea-skim degraded homing in clutter and multipath. Across a terminal period of order ninety seconds the twenty-four-cell launcher is magazine-bound rather than time-bound, servicing twelve tracks at two interceptors each
Close-in gun effective kills per warship	7	short engagement window, low per-shot kill. Ammunition-bound inside the terminal period at of order eight engagements off a ready drum, with no reload available in the time, and a depression-angle and safety-arc limit on fire across the ship's own deck

Parameter	Value	Basis
Directed-energy effective kills per warship	3	hard kill only. A beam is a point effect, so a kill is burn-through on one spot of centimetres and dwell rises with the square of the effective spot: against a weaving, rolling body at engagement range that spot is of order a decimetre and dwell is tens of seconds. Two mounts inside a terminal period of order ninety seconds, discounted for marine haze, spray and rain. Seeker dazzle is separate and faster, and is a divert counter in Section 4.6
Combat air patrol kills (carrier group only)	22	ASSUMPTION, set in the defender's favour. A patrol of four to six fighters carrying four to eight air-to-air rounds each has 16 to 48 rounds available, one pass at a simultaneous arrival, and no time to rearm. Twenty-two kills requires most of them to connect against a low, subsonic, low-observable target in look-down.
Soft-kill and electronic-warfare divert	0	a passive EO/IR seeker is not diverted by radio-frequency measures; optical counters are treated in Section 5
Terminal hit probability	0.80	machine-vision aimpoint selection with a 250 kg warhead, which is far less sensitive to terminal miss distance than a light one
Flight reliability (launch to target area)	0.85	mature production line, Tomahawk-class round
Non-decisive aimpoint allowance	1 hit per warship	opportunity hits on aimpoints that present on the closing approach; scales with group size (see below)
Decisive hits per warship to destroy	4	ASSUMPTION
Probability a hit is decisive	0.80	ASSUMPTION; scales with warhead mass and with delayed fuzing
Carrier mission-kill hits	8	ASSUMPTION. Machinery and generator spaces, aviation fuel and ordnance, the hangar deck and its services, the firefighting main. Flight operations end when the ship cannot make wind over the deck, power its catapults and arresting gear, fuel and arm its aircraft, or isolate a fire.
Oiler destruction hits	2	10,500 t of fuel, unarmoured, no effective terminal defence
Unit cost per round	US$300K	build-to-cost target on a GOCO basis

The non-decisive aimpoint allowance. Some hits land on aimpoints that present on the closing approach (sensors, directors, exposed launchers) and damage the ship without ending its fight. The model reserves one such hit per warship, so a larger group carries a larger allowance. It is a modelling margin, not a claim: a machine-vision seeker aimed at critical spaces wastes fewer hits than this, so the real requirement is lower again.

4.1.2 Threat cases

Table 4.2 *The three threat cases: the defence available, and the rounds required to exhaust it and to destroy the group.*

	A. Small SAG	B. Large SAG	C. Carrier group
Composition	1 Type 055, 1 Type 054B, 1 oiler	1 Type 055, 1 Type 052D, 2 Type 054B, 1 oiler	1 carrier, 1 Type 055, 1 Type 052D, 3 Type 054B, 1 oiler
Area SAM cells	94	156	175
Point-defence cells	48	96	144
Maximum kills the defence can achieve	90	163	237
Rounds to exhaust the defence	105	192	279
Hits required to destroy the group	14	26	41
Rounds to destroy the group	126	230	339
Containers required (9 rounds each)	14	26	38
Batteries required (144 rounds each)	1	2	3
Group replacement value	US$1.45bn	US$2.40bn	US$7.55bn
Cost exchange	38 : 1	35 : 1	74 : 1

The small-SAG composition is the shape observed in the Task Group 107 deployment. The carrier screen (one Type 055, one Type 052D, three anti-submarine frigates and a replenishment ship) is the shape observed escorting the Liaoning group, sized against the submarine threat a carrier most fears rather than for maximum air defence, so the frigates carry a mixed loadout and the group's area cell count is lower than its ship count would suggest. Group replacement values are hull costs plus, for the carrier, its air wing.

The two columns are thresholds, not selectable effects. How many batteries to fire is decided at firing rather than at procurement, but the effect is not thereby calibrated. Exhaust and destroy bound the requirement in the model; the defender's actual loadout, readiness and performance on the day are unknown, so a raid sized to exhaust may destroy, and one sized to destroy may not. Neither column offers a non-destructive option.

4.1.3 Sensitivity

The principal assumption is the Type 055 SAM fraction, because it multiplies through the whole of Table 4.2:

Table 4.3 *Sensitivity to the Type 055 area-cell assumption: the battery ladder holds across the realistic range.*

Type 055 area cells assumed	Small SAG rounds	Large SAG rounds	Carrier group rounds	Battery ladder
50	111	215	325	1 / 2 / 3
60	117	221	331	1 / 2 / 3
75 (baseline)	**126**	**230**	**339**	**1 / 2 / 3**
90	135	239	348	1 / 2 / 3
100	141	245	354	1 / 2 / 3
112 (all cells)	148	252	361	2 / 2 / 3

The ladder holds across the realistic range of loadouts. Only at the physically impossible extreme of all 112 cells loaded as long-range interceptors, with no anti-ship or land-attack missiles aboard, does the small-SAG case tip to a second battery. That is the useful property: the force structure does not depend on getting the loadout assumption right within any realistic range.

The attacker-side parameters matter less than intuition suggests, because the magazine term dominates:

Table 4.4 *Sensitivity to the attacker-side parameters, terminal hit probability and flight reliability, against the large surface action group.*

Terminal hit probability	Large SAG rounds		Flight reliability	Large SAG rounds
0.60	243		0.73	268
0.68	237		0.80	245
0.72	234		**0.85 (baseline)**	**230**
0.80 (baseline)	**230**		0.90	217
0.88	227			

A forty-seven per cent swing in terminal hit probability moves the raid by seven per cent. This is the structural point restated: the early rounds are spent emptying the magazine whatever their accuracy, so accuracy enters only at the margin. It is also why this section does not rest its case on the seeker performing as advertised.

4.2 Rounds needed, and what one battery does

On the single-salvo, sea-skimming profile, against the three cases:

1. **Small SAG** (one cruiser, one frigate, one oiler): 105 rounds to exhaust the magazine, 126 to destroy the group. One battery.
2. **Large SAG** (one cruiser, one destroyer, two frigates, one oiler): 192 to exhaust, 230 to destroy. Two batteries.
3. **Carrier group** (carrier, one cruiser, one destroyer, three frigates, one oiler): 279 to exhaust, 339 to destroy. Three batteries.

A single ship is a different problem from a group, and worth pricing separately. A group is beaten by exhausting its shared magazine (Section 4.1.2); a single ship has none to pool, cannot reload, and against a coordinated sea-skimming salvo is saturated well below the interceptors it carries. On that basis, as rounds and as a fraction of one battery:

Table 4.5 *Rounds to destroy a single ship, and the fraction of one 144-round battery.*

Maritime Targets	Rounds to destroy	Fraction of battery (144-rounds)
Type 055 cruiser	37	0.26
Type 052D destroyer	30	0.21
Type 054B frigate	21	0.15
Type 075 amphibious assault ship	14	0.10
Type 071 amphibious transport	11	0.08
Replenishment oiler or auxiliary	4	0.03

The single-ship basis. A lone ship cannot bring its long-range magazine to bear against a sea-skimming salvo: the missiles cross its radar horizon only on the final approach, so the long-range cells never engage. Its defence is its medium-range air-defence layer, and that is bounded not by how many missiles it holds but by how many engagements it can run at once in the short window before impact. Benchmarked against allied active-homing systems, that fire-control capacity is of the order of sixteen simultaneous engagements for a Type 055, twelve for a Type 052D and eight for a Type 054B; the amphibious ships carry point defence only, and the auxiliary almost none. With two interceptors committed to each incoming round and little time to re-engage, that capacity (not the magazine) sets how many the ship kills before it is saturated. Decisive hits are scaled by hull size, the amphibious ships are counted as mission-killed rather than sunk, and terminal hit probability and flight reliability are as in the group model. The figures are set on the defender's side, crediting the maximum plausible number of those engagements, and they sit far below the group cases because a single ship has no shared, unreloadable magazine to force the attacker to exhaust.

The raid sizes are minimums. Under-resourcing does not fail proportionally: a raid that stops short of exhaustion leaves a defended group afloat and alerted. And if the model errs in the defender's favour, the answer is another battery, which is why the force carries reserve batteries and the ladder is quoted in whole batteries rather than rounds.

4.3 Battery and force structure

The structure builds from the container, with dispersal designed in at every level so no single counterstrike removes a meaningful fraction of the force.

Why sixteen containers. The battery is sized so that the three threat cases fall on clean multiples, and so that they keep falling on them across the range of loadout assumptions a sceptical reviewer might impose:

Table 4.6 *Battery sizing against the three threat cases, and the force structure that follows.*

Containers per battery	Rounds	Small SAG	Large SAG	Carrier group
8	72	2	3	5
10	90	2	3	4
12	108	2	2	3
14	126	1	2	3
16	**144**	**1**	**2**	**3**
20	180	1	2	2

Element	Rounds	Notes
Launch container	9	one 40-foot high-cube, 3 by 3
Battery	144	sixteen containers
Forward force	2,880	20 batteries, four per sector across five sectors
Reserve	720	5 batteries; reloads, reconstitution, partner element
Total	3,600	400 containers; US$1.08 billion fly-away (A$1.81 billion)

The battery is the unit of effect. One battery destroys a small SAG outright, or exhausts a large one. Two destroy a large group. Three destroy a carrier group. Those are the numbers a planner needs and they are integers, which is the property the battery size was chosen to give them.

What the force can fight. Twenty-five batteries is twenty-five small-group engagements, or twelve large-group engagements, or eight carrier-group engagements, before any reload. The line behind the magazine is what turns that into a campaign.

Why twenty-five batteries, and where. Geography sets the deployed force: five sectors cover the critical approaches at four batteries each, with five batteries in reserve for reloads, reconstitution and a partner element. Containers are set apart so that no single counterstrike removes a battery and no barrage removes a sector.

4.4 Basing, manning and site preparation

Batteries sit on Commonwealth land. The Defence estate supplies a security perimeter the battery inherits rather than generates, tenure that does not lapse, and Defence's own works approvals in place of eight state and territory planning systems. It also supplies the separation explosive ordnance requires: at some 36 tonnes of net explosive quantity per battery, the set-back distances from inhabited buildings and public roads run to hundreds of metres, which is the spacing dispersal wants in any case. Where the estate does not reach a sector, land is acquired under the *Lands Acquisition Act 1989*, though acquisition is gazetted and runs one to two years, so it is the exception. No single site is load-bearing: moving a launcher 50 km changes a 2,000 km arc by about two and a half per cent.

The unit is an Army artillery battery. The position holds explosive ordnance, cryptographic material and mission software, so those who guard it, load it and account for it must be cleared and under command, and land fires are the artillery's business. The battery provides security and surveillance, maintenance and monitoring of the containers and the rounds in them, transport, and the crew-mediated software load described in Section 5.4. That is manning in the tens a battery and under a thousand across the enterprise, in trades the Army already recruits. Fire missions are directed centrally through the network of Section 2; the battery is Army-raised and joint-tasked.

Guarding is by monitoring rather than by sentries. A battery is sixteen armed containers and between thirty-two and forty-eight decoys, at spacings the local geography sets, and the batteries of a sector may sit hundreds of kilometres apart. No manning model posts a sentry on each, so positions are instrumented and a reaction element is held centrally. Patrol and maintenance cover decoys identically to live positions, because a container that is never visited becomes identifiable within weeks.

Site works are decided site by site. Each position needs a prepared and drained hardstand, lightning protection and a fire break, and each decoy position needs the same. All-weather access is the expensive question, and it is a separate one: deployment is static, so a battery that cannot be reached in the wet can still fire, and what it loses is reload and physical maintenance until the roads dry. Where new access must be built the line is of order A$10 to 30 million a battery, so the choice at each site is between building the road and holding rounds and spares forward before the wet. The work sits inside the Integrated Investment Program's northern-bases and estate envelope.

Approvals set the schedule, not construction. Environmental assessment of works on Commonwealth land, Indigenous heritage and native-title processes over much of the northern estate, and explosives licensing of each position all run in years and are sequential with survey. Site selection should begin with the development programme rather than with the first delivery.

4.5 Cost-exchange, reload asymmetry, and the line

The cost-exchange is decisive. One battery is 144 rounds at US$300,000, some US$43 million. Against a small SAG worth US$1.45 billion, that is thirty-eight to one in Australia's favour. Against a carrier group, seventy-four to one. These ratios are before counting that the fires are sovereign and resuppliable while the ships are not replaceable in any relevant timeframe. The attack is affordable even when inefficient, which is the property that makes the deterrent credible: the adversary must believe Australia can do this repeatedly, and the arithmetic says it can.

Reload asymmetry is the argument that closes the deterrence case. A vertical launch system cannot be reloaded at sea. The consequence is severe:

A group that successfully defends against the raid has still lost the deployment. It has fired the part of its magazine that engages what Australia can send again, and a group holding full long-range cells but empty medium-range and point-defence cells cannot face a second salvo. It cannot refill without returning to a port, and the ports that can service it are themselves within reach of the same rounds. A raid that is defeated still achieves the operational effect.

This inverts the usual burden. The attacker's failure case is a delay and US$43 million. The defender's success case is an empty magazine, a long transit, and a station that is now uncovered. There is no symmetric argument available to the defender, because there is no way to defend that does not consume the magazine.

The interceptor exchange, before any ship is hit. A large SAG carries 156 area cells and 96 point-defence rounds. Public unit costs for the HHQ-9B and HHQ-16 are not available; the nearest Western comparators, the SM-2 and SM-6, run at roughly US$2.1 million and US$3.3 million each. Even assuming the Chinese equivalents cost a fraction of that, the interceptors expended in defeating one raid cost several times the raid. The exchange is adverse to the defender at the interceptor level alone, before the deployment, the ships or the political effect are counted.

Cost imposition on the adversary. Read from the other side: each escort added to a group raises the raid required against it by roughly 45 rounds for a frigate and more for a heavier ship, some US$14 to 23 million, against the several hundred million dollars the escort costs to build, crew and sustain: a cost-imposition ratio running twenty-to-forty-fold against the defender before a shot is fired.

Unit cost and the production basis. The US$300,000 figure is a build-to-cost target on a GOCO basis, not an extrapolation from Western prime pricing: the state owns the plant, the design authority and the intellectual property; a commercial operator runs production; the engine is additively manufactured; and the guidance electronics are the commodity content Section 5 prices. It is the production model proven at war in Ukraine and now adopted by the United States' own affordable-mass programmes at the same class of unit cost. The sensitivity is linear and bounded: at US$400,000 a round, raid costs rise by a third and every exchange ratio in this section survives.

Rate is the metric behind the inventory. The 3,600-round magazine is a buffer; the asset is a warm line tooled to rebuild the full magazine in about a year at surge. Because capability lives in software, banked rounds take the line's upgrades rather than ageing in their containers, so the war-stock gains capability with age rather than losing it.

Vertical integration to hold the unit cost. The deterrent rests on the unit cost, and the cost is held by owning the highest-value content (the seeker, the autonomy stack and the engine) rather than buying it through layered primes; the acquisition structure that preserves this is Section 5's subject.

4.6 Counters and caveats

The central claim is untested at scale. No massed cruise-missile raid has yet been flown against an alerted naval group. The largest real engagements, some thirty weapons arriving in a trickle, sit well below the threshold this section describes, and they were defeated. The claim rests instead on the defender's own magazine arithmetic, on the land-war evidence that massed precision defeats defended targets past the servicing rate, and on the LRASM's validation of the autonomy stack at the exquisite end of the cost spectrum [1]. The untested step is the conjunction at scale against a warship, and the force treats it as untested: over-matched sizing, a magazine deep enough to be wrong once, and terminal parameters measured rather than assumed.

Every counter converts to raid size. Each falls into one of three families, and each works by raising the size of the raid required rather than by closing the exchange.

Table 4.7 *Counter families, and the design answer to each.*

Counter family	Examples	Design answer
Magazine limited	Deeper interceptor inventory; a larger escort screen	Raid size, priced at Table 4.2. A deeper magazine is exactly what the model counts, and Section 4.1.3 shows the ladder holds to implausible loadouts.
Rate-limited	Close-in guns; hard-kill directed energy	Arrival density in the terminal minute. A defence that can service only so many targets a minute is beaten by how many arrive inside that minute, not by the raid's total size. Marine atmosphere degrades a beam at sea-skim altitude, and the missile's structure lengthens dwell-per-kill. A beam is a point effect, so a rolling and weaving body is a far harder target than the drones these systems were sized against. Directed energy improves on the adversary's refit tempo; a counter-adaptation reaches the whole magazine by software upload in days.
Divert-based	Electronic warfare; multispectral obscurants, decoys, dazzle	The model credits these with nothing: a passive EO/IR seeker is not diverted by radio-frequency measures. Against optical counters the answers are seeker-exposure discipline, aimpoint-memory inertial coast, multi-mode seeking, and the continuous retraining a sovereign line supports.

Where the defender's difficulties compound. The model treats the defence as an orderly sequence, which is generous. In the terminal minute the defender must detect early without gaining lethality, spend two interceptors per track with no time to assess, tell kills from debris, fight on as hits eat sensors, take decisions arriving faster than a crew can process them, and hold a single fused picture across ships whose antennas are being struck [2]. These difficulties do not add; they compound. Combat-management overload is a second exhaustion point that can arrive before the magazine empties, and the model claims nothing for it: it is deliberately pessimistic about the weapon it models. The asymmetry is structural, and it is carried mainly by the ship.

What would falsify this section. Three findings would materially change the conclusion, and it is worth naming them so a reader can watch for them. A demonstrated at-sea vertical-launch reload capability would remove the reload asymmetry. A directed-energy weapon with an effectively unlimited magazine and a kill rate above the arrival rate, in marine atmosphere at sea-skim altitude, would break the rate-limited family. An area interceptor whose per-shot kill probability against a wave-top skimmer approached its performance against high-altitude targets would make the magazine effectively deeper. None is in evidence. All three are worth watching, and the force structure is deep enough to absorb one of them.

References

All references were accessed during preparation of this volume in July 2026. Several of the systems cited are in rapid development; a reader revisiting later should re-verify, as figures move. Quantitative values in the effectiveness model are illustrative planning estimates offered to structure the analysis, not measurements, and the systems below are cited to establish that the LRSM's component choices are fielded or in production, not that any of them is the proposed design.

[1] Lockheed Martin AGM-158C LRASM product literature; USNI News and US Navy budget-justification documents. Cited for the fielded validation of autonomous routing, passive multi-mode terminal seeking and cooperative target allocation in an anti-ship weapon, at the exquisite end of the cost spectrum whose affordable end the LRSM occupies.

[2] The Cooperative Engagement Capability, Johns Hopkins APL Technical Digest 16, no. 4 (1995): 377–396, the developers' own account of a networked fire-control system required to hold 'an identical picture at each unit, of sufficient quality to be treated as local data for engagements', dependent on force-wide gridlock, jam-resistant communications and positive identification. With 'Distributed Multisensor Data Fusion under Unknown Correlation and Data Inconsistency', Sensors 17, no. 11 (2017): 2472, on the intrinsic fragility of a shared track picture under cross-sensor correlation, data incest and clutter. Cited for the difficulty of sustaining cooperative engagement under saturation and degradation.

[3] Wikipedia, 'Type 055 destroyer'; Naval Technology, 'Type 055 Class Destroyers'. Cited for the defended-group air-defence picture: 112 universal VLS cells (64 forward, 48 aft) carrying HHQ-9B long-range surface-to-air missiles, a 24-cell HHQ-10 point-defence launcher and a 30 mm close-in system. The surface-to-air fraction of the mixed loadout is the model's principal assumption and is swept in Section 4.1.3. The HHQ-9B is an active radar homing missile and the class carries no separate illuminators, which is why no fire-control channel limit is assumed anywhere in this section.

Section 5 — Industrial Partnerships and Sovereign Production

This section maps each key capability to a realistic production path (sovereign design-and-production, joint design with sovereign manufacture, or international collaboration where appropriate) and names candidate industrial partners. Its purpose is credibility: to show that the industrial base to build this architecture largely exists and is identifiable, and that much of it is already standing up under the Guided Weapons and Explosive Ordnance (GWEO) Enterprise [1] and the autonomous-systems programs. Partners are named as illustrative of the available base and current arrangements as of mid-2026, not as asserted commitments to this architecture; figures are drawn from public reporting.

5.1 The principle: sovereignty graded by capability

Sovereignty is not all-or-nothing; it is graded by how much of the design, the intellectual property, and the production Australia holds. The architecture applies the most self-reliant path each capability allows, which is the strategy's own test of where self-reliance is feasible and advisable: full sovereign design-and-production where it is achievable and where the capability and IP can be held domestically (the sovereign missile is the strongest case); joint design with sovereign manufacture where partnering de-risks development and secures resupply; and international collaboration or acquisition where the self-reliant path is not realistic, accepting the dependency but diversifying it. The decisive test is resupply without permission in a crisis, which favours pushing each capability as far toward sovereign production as it will go.

The 2026 Integrated Investment Program supplies the frame [2]. Its accelerated priorities include expanding long-range strike, fielding autonomous and uncrewed systems at scale, and, verbatim, delivering “a resilient and secure multi-orbit satellite communications system”; the accompanying commitment is to “build stronger, more diverse international industrial partnerships and supply chains”. Every industrial path in this section sits inside a priority the government has already funded and named. The section is not proposing a new industrial policy; it is directing an existing one.

5.2 The Joint Sovereign Targeting Network (JSTN): the keystone

The JSTN’s industrial question is not sustaining what exists but building what does not. The network’s in-service spine (the Jindalee over-the-horizon radar network, the E-7A Wedgetail fleet, and the joint air and missile defence sensors, with their established Australian sustainment base) is funded and industrially anchored, and Section 2 treats it as given. Three things must be bought or built.

The SAR sensing constellation: purchased, Australian-owned, Australian-operated. The maritime custody layer is bought as a complete national system, owned and operated by Australia, on the model that ICEYE has now delivered to five European governments, at the delivered prices and timelines quantified in Section 2.9 [3]. That is the ownership-and-operation model this architecture specifies. ICEYE is a candidate supplier, no more: the specification is the purchased, Australian-operated constellation, whoever builds it.

The ELINT/datalink constellation: purchased, with jointly developed payloads. The ELINT/datalink shell carries custom payloads (the ELINT receivers and the assured L-band fire-control bearer of Sections 2.4–2.5) that do not exist as catalogue products and would be jointly developed with the vendor, with Australian design authority over the electronic-intelligence payload and the targeting waveform (Section 2.4). Rocket Lab and SpaceX are the candidate partners.

Rocket Lab builds spacecraft, operates its own launch (Electron, with the medium-lift Neutron in development), and its acquisition of Iridium (announced June 2026, expected to close mid-2027 subject to approvals [4]) would bring an operational global L-band constellation, its scarce internationally coordinated spectrum, and replacement launch under one roof: constellation, spectrum and reconstitution from a single partner. SpaceX's Starshield line offers communications and hosted payloads on the Starlink production base, proven at scale for United States national-security customers, though not yet operated sovereignly by any ally [5].

Under either vendor the constellation is bought into Australian ownership and operation, with a sovereign ground segment and sovereign cryptography; the vendor is a supplier and a reconstitution partner, not the owner. That distinction is what makes the shell pass this section's own test.

The sovereign core: the decisive build. The element that must be Australian-held outright is the software: the basestation interfaces into the network, the data-fusion engine, the AI imagery and signals analysis, and the weapon datalink, because find-fix-fire authority must rest in Australian hands.

The build draws on a real domestic base: the Defence Science and Technology Group and the Australian C2 and autonomy firms for the sovereign core, and CSIRO for the processing-and-analytics foundation. CSIRO already operates the NovaSAR-1 SAR satellite tasking share as a national facility, runs a Centre for Earth Observation and a cloud-scale Earth-analytics platform, and leads the AquaWatch sensing mission with the SmartSat CRC [6], which is a sovereign SAR-processing and analytics competency in being, not a hope.

The cost structure is settled policy. Defence has already ruled on how Australian satellite capability should be bought. JP 9102, the exquisite geostationary SATCOM program for which Lockheed Martin was selected in 2023, commonly reported at about A$7 billion, was cancelled in November 2024 on the department's own stated ground that a single-orbit GEO system "would not meet strategic priorities" and that "Defence must instead prioritise a multi-orbit capability increasing resilience" [7]; the 2026 IIP then made the multi-orbit system an accelerated priority [2].

The commercial proliferated-LEO constellations (ICEYE, Iridium, Starlink) exhibit the cost structure that makes this affordable: commercial-derived satellites on a shared component base, priced at marginal cost rather than as bespoke prime-vendor developments (Glossary, GOCO; the delivered-price benchmarks are in Section 2.9). The JSTN's space layer buys the way Defence has already said satellite capability must be bought.

5.3 Long-Range Precision Fires: the industrial base

The GWEO Enterprise gives the architecture's fires a funded industrial setting. Australia has committed up to A$21 billion (rising toward A$26–36 billion in the 2026 IIP) to Australian guided-weapons manufacturing, with three strategic partners: Lockheed Martin Australia, Raytheon Australia and Kongsberg Defence Australia [1]. The architecture's retained tiers are already in production or construction under it.

Table 5.1 *Long-range precision fires: production path, partner and status by capability.*

Capability	Path & partner	Status (mid-2026)
GMLRS (short-range fires)	Joint design, sovereign manufacture: Lockheed Martin Australia	In production at Port Wakefield, SA since December 2025; the first rounds made outside the US
PrSM (precision strike)	Co-development / co-production: Lockheed Martin, US MoU	Inc 1 (~500 km) in service; Inc 2 (LBASM), the anti-ship increment co-developed with Australia, first flight test 2026; Inc 4 ~1,000 km; Australian assembly optioned
NSM / JSM (anti-ship / strike)	Joint, sovereign manufacture: Kongsberg Defence Australia	Williamtown (Newcastle) factory in construction; production from 2027, the first outside Norway
Hypersonics (premium end)	Co-development: Raytheon/Northrop (HACM) and allied programs	The premium tier remains allied co-development (Section 1)

GWEO plans a new Australian weapons manufacturing complex for high-rate production: reporting points to a goal of up to ~4,000 missiles a year by the end of the decade, an order of magnitude beyond ADF needs, with the surplus exported into the allied supply chain [1]. That export logic is what makes the GMLRS line economically viable at rate. A solid-rocket-motor manufacturing complex (Mulwala) and explosive-assembly facilities (Orchard Hills) are funded to close the component gaps.

What this base supplies to the sovereign round is specific: the settled precedent that guided weapons are manufactured in Australia with state capital; the energetics and explosive-assembly infrastructure at Mulwala and Orchard Hills that any Australian-built round draws on; and production of the complementary tiers the architecture retains. What it does not supply is the price class. Each existing line manufactures a foreign prime's design under licence, with design authority remaining abroad and unit costs in the band priced in Section 5.4. These are valuable lines, and licensed production is the right way to buy their tiers; they are also, organisationally, the traditional model, and no extension of them produces a US$300,000 round.

What Section 5.4 takes up is therefore not an extension of the existing lines but the piece they do not supply: the sovereign missile, and the production organisation that determines its price.

5.4 The sovereign missile: a mature class, produced differently

The missile is where full sovereign design and production is most achievable and most valuable, because the airframe is conventional, the intellectual property can be held domestically, and resupply without permission is the whole point of Australian mass. The domestic base is real and growing: the 2026 IIP commits up to A$22 billion to drone, counter-drone and autonomous-system technologies, with at least A$4.3 billion in uncrewed aerial systems [2]. This section makes four arguments in sequence: the munition class is mature; the requirement that remains is price, which is a production problem; the production organisation is what sets the price; and two credible paths exist to that organisation, each with trade-offs the Government can weigh.

The class is mature

A 1,500–3,000 km conventional cruise missile is five-decade-old technology that many states now field [13]:

Table 5.2 *Fielded long-range conventional cruise missiles.*

Country	Missile	Launch weight	Range	Fielded
United States	Tomahawk Block IV/V	~1,600 kg	~1,600 km	2004 / 2021
Russia	Kalibr 3M14	~1,700–2,300 kg	~1,500–2,500 km	~2015
China	CJ-10	~1,500–2,000 kg (est.)	~1,500–2,000 km	mid-2000s (est.)
France	MdCN	~1,400 kg	>1,000 km	2017
South Korea	Hyunmoo-3C	~1,360 kg	~1,500 km	2012

Air-launched siblings (Kh-101, CJ-20) extend the same designs, Iran and North Korea field or claim the class, and Pakistan's Babur. The class is within reach of any serious industrial economy, and Ukraine has just fielded it in three years under bombardment. What is genuinely new in this architecture is not the missile; it is producing the missile at Table 5.4's price, which is a production-model problem, not a physics one.

Representative example and candidate collaboration: Destinus Ruta. That the long-range rung is reachable on a modern production model is shown by the European firm Destinus, whose Ruta family is a turbojet-powered, EO/IR-AI-guided, GNSS-denied cruise system [19]. Block 1 has been used in combat in Ukraine; Block 2, container-launched with a 250 kg warhead beyond 700 km, flew in April 2026 and has 700 rounds on order for Ukraine; Block 3 extends the same stack past 2,000 km and is in accelerated development with Rheinmetall, whose joint venture is scaling production on an all-European supply chain, the same independence logic Australia seeks.

Ruta is cited as a representative example of the class and a candidate collaboration or technology-transfer partner, not as an airframe to licence: the LRSM is an Australian design. With the fielded systems at Table 5.2, the point stands that the missile proposed here is a class that exists rather than a class that has to be invented.

The requirement is price

The two production models, priced. The unit cost this architecture rests on is not a discount on the prevailing price; it is what a different production model produces. Both ends of the comparison are now observable in programmes of record.

Table 5.3 *Strike weapons on the prime production model: flyaway unit costs.*

Strike weapon (prime model)	Unit cost
Tomahawk Block V	~US$1.9M
Naval Strike Missile	~US$1.9M
LRASM	~US$3.2M
Precision Strike Missile Inc 1	~US$2.0M

These are flyaway figures, drawn from United States budget justification documents [8], and they are the conservative choice, excluding canisters, spares, support equipment and non-recurring costs. The appropriated lines run roughly double: the FY2027 request buys 785 Tomahawks for US$3 billion, some US$3.8 million a round [9]. Foreign military sales, which is how Australia buys, land in the same higher band once canisters, integration and sustainment are bundled. The lower figures are quoted throughout because they make the gap argued here smaller and are the harder position to attack.

Table 5.4 *The volume-manufacture model: unit-cost targets and programme scale.*

Volume-manufacture model	Unit cost	Scale
Family of Affordable Mass Missiles (USAF)	~US$220,000 (reported)	28,000 rounds over five years
Low-Cost Containerized Missiles (DoD)	fixed-price; target undisclosed	10,000+ rounds from FY2027
CHAOS (USN)	undisclosed	prototype, for coalition sale
FP-5 Flamingo (Fire Point, Ukraine)	~US$500,000–600,000 (reported)	combat-used from August 2025; ~3,000 km, 1,150 kg warhead
Sovereign round (this architecture)	US$300,000	3,600 rounds

The United States rows are funded programmes built on commercial vertical integration, additive manufacture at rate, commodity sensing and compute, and design for manufacturability from the outset [10]. The Ukrainian row is the wartime existence proof: a mid-size economy under bombardment fielded a 3,000 km cruise missile at roughly a sixth of a Tomahawk's price within three years of the company's founding [11]. Its figures are wartime company claims made under an active Ukrainian procurement investigation, and they are cited here for the demonstrated capability and price class, not for the firm's production-rate claims; the strikes themselves, at ranges out to ~1,400 km, are documented.

The gap between Tables 5.3 and 5.4 approaches an order of magnitude, and it is a gap in method rather than in physics: the guidance content that once justified a seven-figure price is now commodity, and what remains expensive is engine, airframe and militarised qualification. The ~US$300,000 round is priced against the second table, and the production practices and organisation that keep it there are the subject of the remainder of this section.

Production is where the innovation lives

Because the class is mature, the innovation this architecture requires sits almost entirely in production engineering. Four practices, each demonstrated in the Table 5.4 programmes, set the price [10][12].

The engine is additively manufactured. The turbojet has been the historical long pole in cheap cruise-missile production, and additive manufacture has moved it onto a falling cost curve: Anduril developed Roadrunner's turbojet in-house specifically to control cost, with unit cost falling with rate [12]; Destinus passed its thousandth in-house turbojet in June 2026, serial cruise-missile propulsion built inside Europe with design, tooling and testing held internally; and Pratt & Whitney has ground-tested its TJ150 expendable engine with roughly sixty per cent of its hardware additively manufactured [20]. The production engine is additively manufactured in Australia from the first serial block. Its design is either purchased outright, which buys schedule, or developed internally on the path Anduril has demonstrated, whichever closes faster against the production date. A purchased design is bought, not rented: it transfers with no per-round levy on domestic production, with any residual licence fees confined to export sales. Prototypes fly on purchased engines; at the volume line's cost target an imported engine carries a supplier's margin and a foreign serial dependency the missile exists to avoid.

Components are commercial wherever qualification allows. The guidance, navigation and compute content is commodity: commercial-grade sensing, processing and inertial components qualified for the round's short one-way flight life rather than for decades of shipboard storage. Militarised qualification is confined to the elements that need it (warhead, safe-and-arm, the hardened datalink), because qualification cost, not capability, is a large part of what separates Table 5.3 from Table 5.4.

The airframe is tooled for volume. Die-cast or extruded aluminium structural sections, injection-moulded polymer and composite components, CNC milling confined to interfaces and precision fits, robotic joining, and automated inspection and test: designed for manufacturability from the outset rather than converted from a hand-built prototype. The consolidation practice is the one automotive manufacture has demonstrated, single castings replacing fabricated assemblies of dozens of parts, applied here at missile scale and conventional press tonnage. The airframe is deliberately conventional so that its cost sits in tooling amortised across the run, not in touch labour.

The line is highly automated. Automation takes labour cost out of the round, removes the defect modes of manual assembly, and makes rate responsive: production scales with demand by adding shifts and cells rather than by recruiting and certifying a workforce. Responsiveness is the war-stock argument in industrial form: the line, and the demonstrated rate at which it can surge, is as much a part of the deterrent as the rounds it has already delivered.

Serial production proceeds in blocks in the usual military sense, each block a fielded hardware configuration. The initial production block is sovereign in the airframe, the integration, the seeker, the autonomy, the software, the launch system and the JSTN integration, with engines Australian-built from the first serial round. Later blocks group hardware increments in range, seeker and engine refinement.

Software is not block-bound. The missile carries an open software architecture, and software updates, including machine-learning targeting models, are distributed fleet-wide and loaded by battery crews from secure hardware at the container. The update path is physical and crew-mediated by design, keeping software loading off the network attack surface; the JSTN carries cues, tracks and fire orders, not code.

The staging extends an accepted trajectory: it is the pattern the enterprise already follows, from licensed GMLRS assembly at Port Wakefield toward progressively deeper domestic content. The initial block remains a genuine design programme with schedule risk of its own, not a by-product of an existing line; a sovereign 2,400 km missile with a machine-vision seeker is a low-risk build on a demonstrated technology base rather than a free one.

The organisation is the cost lever

Table 5.4's prices are produced by a kind of organisation before they are produced by a factory. Anduril took Roadrunner from concept to combat evaluation in under two years with an engine built in-house [12]; Perennial Autonomy went from founding to a US$500 million Pentagon contract in about three, on a ~US$5,000 AI-guided strike weapon co-produced with Ukraine [14]; Fire Point went from founding to a fielded 3,000 km missile in about three years of war [11]; SpaceX and Rocket Lab made the same demonstration against the traditional launch integrators such as Lockheed, Boeing and then ULA.

The common properties are cultural and structural, not proprietary: an organisation aligned tightly on the mission, fast learning loops with tolerated failure, engineers rewarded for shipped capability, adequate capitalisation to iterate, and manufacture brought in-house where feasible, which cuts external dependencies, strips supplier margin stacking, and collapses design-iteration time.

What the Commonwealth must hold

The sovereignty requirement is defined by function, not by the manufacturer's share register. Five things must be in Australian hands for the missile to pass this section's test: decision and release authority; the targeting software and the JSTN integration, communications and cryptography around it, because these carry rules-of-engagement weight; the type design and technical data package for the round, held from the outset or transferring at a defined maturity gate; the production line itself on Australian soil, with plant and tooling under Commonwealth title; and the ability to operate and resupply the line if the commercial relationship ends, which means reversion rights written into the capital, not goodwill. An operator's ownership matters where it would put one of these five outside Australian control; it is not, by itself, the test. The state-title template is one Australia already runs at Mulwala and Benalla for propellants, strengthened here by holding the intellectual property as well as the buildings.

Two production options

Two credible paths lead to an organisation that can produce Table 5.4's prices under the holdings above. Both are government-owned, contractor-operated in structure (Glossary, GOCO); they differ in who holds airframe type design at the start and how it transfers. Range and payload are specifications common to both, not discriminators between them.

Option A: a new Australian-owned production organisation, capitalised by the Commonwealth. The state holds the plant, the design authority and the intellectual property from day one; a privately held Australian operator runs the line under a compact that excludes foreign acquisition; the culture is built new, without a legacy integrator's inheritance. This is the cleanest pass of the sovereignty test, and it is the slowest path. The organisation starts with no background intellectual property and no standing engineering team, and the culture that produces volume prices cannot be specified into existence; it follows from the contract's incentives and from the people recruited, and both take time. The Ghost Shark precedent must be read carefully here: its speed from concept to production contract rested on background intellectual property the contractor brought from prior work, and only its contracting tempo, not its development tempo, transfers to an organisation starting cold [12].

Option B: Anduril Australia engaged to manufacture. The Commonwealth contracts an established volume manufacturer whose production model and culture already exist, on terms structured in two instruments: a cost-shared development contract to design freeze, then fixed-price production per round and per container, plus an annual capacity fee that holds the line warm at a demonstrated surge rate. Commonwealth capital into the plant and research is structured as asset title with reversion, not as a grant. Type design and the technical data package transfer to the Commonwealth at a defined maturity gate, with the contractor remaining design agent; the software interface, JSTN integration, communications and cryptography are Australian throughout. Third-country sales are a matter between governments, which concedes a practical United States veto; the capacity fee is therefore a permanent line item in the costing, not a cost that exports amortise away. What Option B buys is time and risk: an in-house additively manufactured turbojet already fielded, a demonstrated production culture, and contracting tempo shown in Australia on Ghost Shark, whose co-development contract was placed in 152 days against a usual Defence cycle of one to three years [12]. What it costs is a manufacturer whose parent is United States-owned, with the export-control gravity that brings, and a predictable negotiation over the software layer, where the contractor's commercial value sits. The architecture limits that exposure structurally: the round's intelligence lives in the JSTN, not in the missile, so the software at issue is small and the interface specification, which the Commonwealth writes, is the control point.

Table 5.5 *The two production options compared.*

	Option A: new Australian organisation	Option B: Anduril Australia manufacture
Airframe type design & TDP	Australian-held from the outset	Transfers at a maturity gate; contractor as design agent
Sovereign core (targeting software, JSTN, cryptography)	Australian	Australian
Plant and tooling	Commonwealth title	Commonwealth title with reversion
Time to first production block	Longest: team, IP and culture built from zero	Shorter: existing model, engine and culture
Production culture	Must be created by contract incentives	Demonstrated, imported with the operator

	Option A: new Australian organisation	Option B: Anduril Australia manufacture
Engine path	Design purchased or developed; AM line stood up new	In-house AM turbojet precedent (Roadrunner)
Export path	Australian decision, strictest settings	Between governments; practical US veto
Cost structure	Capital plus operating contract	Capital, fixed price per round, permanent capacity fee
Principal risk	Schedule and organisational formation	Software-layer negotiation; parent ownership

The selection between them is a judgment the Government can make on evidence rather than doctrine: time to the first block at rate, the surge rate demonstrated rather than certified, whole-of-life cost including the capacity fee where one is paid, and the five holdings above, which are the constant. Under either option the organisation is contracted the way this section has argued volume prices require. Contracted the way exquisite platforms are contracted, the price class is forfeited before the first round is built; contracted without Australian control of the five holdings, the self-reliance is forfeited instead.

The sovereign development items

The development items below are the same under either production option, and they are where the sovereign value concentrates.

Table 5.6 *Sovereign development items: path, candidate partner and note.*

Element	Path & candidate partner	Note
Airframe & integration	Sovereign design-and-production: Australian primes/SMEs	Simple airframe; the most domestically-producible element
Engine (turbojet)	Sovereign build via additive manufacturing: design purchased outright or developed internally	Imports for prototyping only; the historical long pole, now on the AM curve; Anduril built the Roadrunner turbojet in-house for cost control [12]
Machine-vision seeker	Sovereign design: Australian autonomy/AI firms	The load-bearing development item; a software advantage
Loitering-munition base	Existing SMEs: DefendTex (D40), Sypaq (Corvo family)	Proven Australian loitering-munition design and production
Autonomy / C2	Sovereign: e.g. Anduril [12] Australia (Lattice-class)	Swarm coordination on the fast software curve

The seeker: an Australian base, externally validated. The machine-vision terminal seeker is the load-bearing development item, and Australia is not starting cold. Electro Optic Systems (EOS) [15] provides a combat-proven engineering base in the adjacent disciplines: precision electro-optic detection and tracking of small, fast, manoeuvring aerial targets in clutter, aided target recognition, and exquisite stabilisation and pointing, competencies that transfer directly to a terminal seeker. For the passive anti-radiation (SEAD) seeker variant, CEA Technologies [16] is the natural RF engineering base: detecting, characterising and homing on a hostile emitter is a radio-frequency sensing-and-processing problem in the heartland of the CEAFAR and CEAMOUNT competency.

That cheap machine-vision terminal guidance works under jamming is validated in combat now: The Fourth Law's onboard module takes over the terminal phase and defeats jamming for a marginal cost increase [17]; Auterion's kits provide combat-proven visual terminal guidance at tens-of-thousands scale, with swarm coordination demonstrated in US live-fire tests [18]; and Perennial's Hornet, now bought by the US Army, does it at ~US$5,000 a round [14]. These are external reference points and possible collaborators, not the sovereign core. The engineering base plus combat-proven validation makes the seeker a realistic Australian build, not a leap.

5.5 Summary: the base largely exists

The credibility point is that this architecture is not a call to build an industrial base from nothing: it is a call to direct one that is already standing up, inside priorities the 2026 IIP has already named. The JSTN's in-service spine is funded; its constellations are purchasable on national-procurement models already delivered in Europe, with credible vendors for the jointly developed payloads; the GWEO Enterprise is delivering the guided-weapons base with named partners and funded facilities; and the autonomy firms and SMEs are delivering Australian mass domestically and fast.

The architecture's distinctive asks are to build the sovereign core, to buy the two constellations into Australian ownership and operation, and to stand up the sovereign missile on an Australian line under one of the two production options in Section 5.4, prioritised as the highest-leverage, fastest-to-field investments.

Three qualifications remain. GWEO is a crawl-walk-run program that critics argue must move faster against the threat timeline. Some elements (Tomahawk, US targeting components) remain imported and should be diversified; South Korea, Britain, Germany and Japan offer credible technology-transfer paths, and ITAR-free European systems another. Named partners reflect current arrangements and capabilities, not commitments to this specific architecture. The industrial base to build the architecture is identifiable and, in large part, already in production.

References

All references were accessed in July–August 2026. Many concern fast-moving commercial and defence programmes and should be re-verified before external use. Named companies are cited for demonstrated capability, not as commitments to this architecture; company ownership, contracts and product specifications change.

[1] Australian Department of Defence, Guided Weapons and Explosive Ordnance (GWEO) Enterprise announcements, 2023–2026: GMLRS production commenced Port Wakefield, SA, December 2025 (the second facility worldwide after Camden, Arkansas); the Australian Weapons Manufacturing Complex sized toward ~4,000 missiles a year by 2029; Rocket Motor Manufacturing Complex planned for Mulwala (Northrop Grumman preferred partner); explosive-assembly facilities at Orchard Hills; strategic partners Lockheed Martin Australia, Raytheon Australia, Kongsberg Defence Australia.

[2] Australian Government, 2026 National Defence Strategy and Integrated Investment Program, media release, 16 April 2026 (accelerated priorities, including “a resilient and secure multi-orbit satellite communications system”; “$425 billion over the decade”; industrial-partnerships commitment); GWEO decade planning of A$26–36 billion and drone/counter-drone/autonomous figures per the IIP and ASPI analysis, April 2026.

[3] National SAR procurement benchmarks: ICEYE POLSARIS (Poland), ~€200m for four satellites plus fixed and mobile ground segment and operator training, contracted May 2025, handed over within twelve months (May 2026); Finnish Defence Forces, ~€158m, September 2025; Royal Netherlands Air Force, Portuguese Air Force and Greek programmes; Rheinmetall–ICEYE contract with Germany, ~€1.7bn over five years, a data-and-operations service rather than a hardware purchase; Swedish Armed Forces, ~SEK 1.3bn across ICEYE and Planet. ICEYE, Rheinmetall and government press releases, 2025–26. Cited for the delivered-system prices, the ownership-and-operation model, and the timelines; derived planning figures are in Section 2.9.

[4] Rocket Lab Corporation SEC filings and press: ‘Rocket Lab to Acquire Iridium’, 29 June 2026 (enterprise value ~US$8bn, expected close mid-2027, subject to shareholder and regulatory approvals); ‘Rocket Lab Completes Mynaric Acquisition’, 14 April 2026 (laser optical communications); Geost (missile-warning and missile-tracking sensors) acquired 2025. A vertically-integrated spacecraft, launch, laser-communications and, prospectively, Iridium-network base; the Neutron medium-lift vehicle remains in development.

[5] SpaceX Starshield: communications and hosted-payload line on the Starlink production base; the NRO’s proliferated architecture (satellites on orbit from 2024, declared operational October 2024) and US Space Force satcom awards; US-contracted Starshield access for Ukraine. US reporting 2023–2026. Operational customers to date are US national-security agencies; no allied sovereign operation yet.

[6] CSIRO: NovaSAR-1 national facility (Australian tasking share of the SSTL SAR satellite, operated as national research infrastructure); Centre for Earth Observation; EASI Earth-analytics platform; AquaWatch Australia mission with the SmartSat CRC. CSIRO publications, 2021–2026. Cited as the sovereign SAR-processing and analytics expertise base.

[7] JP 9102 cancellation: Australian Department of Defence statement, 4 November 2024 (“a single orbit GEO-based satellite communications system would not meet strategic priorities”; “Defence must instead prioritise a multi-orbit capability increasing resilience for the Australian Defence Force”); Lockheed Martin selected as preferred bidder April 2023; program value commonly reported at ~A$7 billion; Deputy Prime Minister remarks, 6 November 2024 (“change of direction”, distributed microsatellite reasoning). Defence and trade press, 2023–2025.

[8] United States budget justification documents, flyaway unit costs: Tomahawk Block V US$1.89M and Naval Strike Missile US$1.90M per The War Zone’s compilation of Navy budget documents (2024); LRASM ~US$3.2M per USAF/USN budget documents (FY2024–25 multiyear); PrSM Increment 1 ~US$2.0M per Army budget documents (varies by year). Flyaway costs exclude canisters, spares, support and non-recurring costs.

[9] USNI News, 6 April 2026, on the FY2027 Navy budget request: 785 Tomahawks for US$3 billion (~US$3.8M a round) and 540 SM-6 for US$4.33 billion (~US$8.0M a round)

across the budget and reconciliation lines. Cited for appropriated per-round figures, which include canisters, spares and support that flyaway costs exclude, and which bracket the foreign-military-sales pricing Australia pays.

[10] United States affordable-mass programmes: USAF Family of Affordable Mass Missiles (framework agreements July 2026; ~28,000 rounds over five years; ~US$218,000 a round reported; awardees include Anduril, CoAspire and Zone 5 Technologies); DoD Low-Cost Containerized Missiles (framework deals May 2026; 10,000+ rounds from 2027; fixed-price, unit target undisclosed); USN CHAOS (CoAspire RAACM-ER, US$70M OTA, July 2026, for coalition partners). Defense News, DefenseScoop, The War Zone, 2026.

[11] Fire Point (Ukraine), FP-5 Flamingo: ~3,000 km, ~1,150 kg warhead, turbofan; unit price ~US$500,000–600,000 as reported (a Czech crowdfunding purchase priced one round at ~US$520,000); combat use from 30 August 2025, with documented strikes including Kapustin Yar and Votkinsk (~1,400 km), February 2026. Reported figures are wartime company claims made under an active Ukrainian anti-corruption investigation into the firm's drone pricing; observed employment is documented, claimed production rates are not. Ukrainian and Western reporting, 2025–2026.

[12] Anduril: Roadrunner / Roadrunner-M reusable VTOL turbojet interceptor, unveiled December 2023, combat evaluation from January 2024, US DoD order for 500+ (October 2024); the turbojet developed in-house for cost control; unit cost in the low hundreds of thousands, falling with rate. Ghost Shark extra-large autonomous undersea vehicle (Anduril Australia): development agreement 2022, A$1.7bn production contract September 2025, concept to production contract in about three years, the co-development contract placed in 152 days (H. Channer, 'Can Defence replicate its Ghost Shark success?', ASPI The Strategist, 17 October 2023). Company reporting and defence press, 2023–2026.

[13] Fielded long-range conventional cruise missiles: US Navy Tomahawk fact file; CSIS Missile Threat entries for 3M14 Kalibr, CJ-10, Hyunmoo-3, Kh-101, Babur; French Ministry of the Armed Forces, MdCN. Figures for Chinese, Iranian, North Korean and Pakistani systems are estimates or national claims, marked accordingly.

[14] Perennial Autonomy (US; formerly White Stork; Ukraine co-production via Swift Beat LLC, July 2025): Hornet AI-guided strike drone, ~US$5,000, optical navigation resistant to GPS jamming; deployed in Ukraine from spring 2025; US$500M three-year US DoD contract, May 2026; adopted by the US Army. US and Ukrainian reporting, 2025–2026.

[15] Electro Optic Systems (EOS) company disclosures and defence press, 2023–2026. Slinger counter-drone remote weapon station (160 ordered for Ukraine, 2023) and the EO/IR precision tracking, aided-target-recognition and stabilisation competencies cited as the adjacent engineering base for terminal-seeker development; part of the successful Leidos team for Project LAND 156 (2025).

[16] CEA Technologies company disclosures and Australian Government / defence-press reporting, 2023–2026. Radar house with Australian-held intellectual property; Australian Government majority ownership (~72%) taken in 2023, announced at ~A$500m; 80%+ Australian industry content; CEAFAR active phased-array radar on Anzac- and Hunter-class frigates; CEAMOUNT illuminator for semi-active homing-missile uplink.

[17] The Fourth Law (Kyiv); reporting 2025–2026. Onboard-AI terminal-guidance module that assumes control in the terminal phase and defeats jamming for a marginal cost increase; combat-used and scaling. Cited as external validation of cheap machine-vision terminal guidance, not the sovereign core.

[18] Auterion; reporting 2025–2026. Terminal-guidance kits (33,000 contracted for Ukraine, 2025) with combat-proven visual terminal guidance; swarm coordination demonstrated in US live-fire testing and being fielded to Ukraine. Cited as external validation and a possible collaborator, not the sovereign core.

[19] Destinus, Ruta product references (Blocks 1–3) and Rheinmetall–Destinus joint-venture announcements, April–June 2026: Block 1 combat-used in Ukraine; Block 2 container-launched, 250 kg-class warhead, >700 km, flight-tested April 2026, 700 ordered for Ukraine by the Netherlands; Block 3, >2,000 km, in accelerated development; production scaling on an all-European supply chain. Company, Rheinmetall and trade reporting, 2026.

[20] Destinus, ‘Engine No. 1000: Destinus reaches industrial-scale turbojet production in Europe’, 18 June 2026: 1,000th in-house T150 turbojet, with design, tooling, test infrastructure, supply chain and quality control held internally; powers Ruta Blocks 1–2, with the T220 planned for Block 3. RTX / Pratt & Whitney, demonstration test of a TJ150 expendable turbojet with approximately 60 per cent of its hardware additively manufactured; company releases and trade reporting, 2025–26. Cited for in-house serial engine production at scale and for additive manufacture of expendable turbojets respectively.