

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'.

*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 used to structure the analysis, not 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. This is his only professional connection to Defence. He writes from a systems-architecture perspective; he is 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.....	4
Section 1 — Joint Architecture Operational Concept.....	6
1.1 Strategic objective.....	6
1.2 Alignment with the Integrated Investment Program.....	6
1.3 Capability overview.....	7
1.4 Scenario one — Task Group 107 custody.....	10
1.5 Scenario two — coercion ladder answered in rungs.....	10
1.6 Scenario three — maritime strike with fires.....	11
1.7 Scenario four — joint maritime strike with hypersonics.....	11
1.8 Scenario five — joint land attack with SEAD.....	12
Section 2 — Joint Sovereign Targeting Network: Requirements and Design Concept 14	
2.1 Purpose and the joint imperative.....	14
2.2 Requirements.....	14
2.3 Design concept: the layers.....	15
2.4 The space layer: two shells, two procurements.....	16
2.5 Imaging reconnaissance and land-attack scene matching.....	19
2.6 The JSTN Core: maritime fusion, engagement management and the weapon datalink.....	20
2.7 Resilience and reconstitution.....	21
2.8 Building on existing investment.....	22
2.9 Sizing and cost: the trade space.....	22
2.10 Development sequence, and the interim.....	24
References.....	25
Section 3 — Long-Range Precision Fires: Requirements and Design Concept.....	29
3.1 Requirements.....	29
3.2 Long Range Strike Missile.....	30
3.3 Launch: the 40-foot high-cube container.....	32
3.4 The attack profile.....	33
References.....	35
Section 4 — Fires Effectiveness and Force Structure.....	36
4.1 The effectiveness model.....	36
4.2 Rounds needed, and what one battery does.....	41
4.3 Battery and force structure.....	42
4.4 Basing, manning and site preparation.....	42
4.5 Cost-exchange, reload asymmetry, and the line.....	44
4.6 Counters and caveats.....	45
References.....	46
Section 5 — Industrial Partnerships and Sovereign Production.....	48
5.1 The principle: sovereignty graded by capability.....	48
5.2 The Joint Sovereign Targeting Network (JSTN): the keystone.....	48
5.3 Long-Range Precision Fires: the industrial base.....	50
5.4 The sovereign missile: a mature class, produced differently.....	51
5.5 Summary: the base largely exists.....	57
References.....	57

Glossary

The architecture

Battery	Sixteen 40-foot high-cube launch containers at nine rounds each: 144 rounds.
JSTN	Joint Sovereign Targeting Network. A persistent, theatre-wide surveillance, reconnaissance and targeting network, optimised for maritime operations across Australia's approaches. It is built around Australian-owned SAR, RF/ELINT and communications constellations. It produces targeting-quality maritime tracks for the architecture's fires and contributes maritime and ELINT data to the wider Defence targeting and battle-management environment, 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.

Surveillance, reconnaissance, 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, classification and geolocation of an emitter. Radar and datalink characteristics can support platform classification for both surface and airborne emitters.
ELINT/ comms shell	The upper of the JSTN's two low-earth-orbit constellations: 180 cross-linked satellites at ~780 km in twelve planes at 45° inclination, carrying passive ELINT receivers able to geolocate and classify surface and airborne emitters, the assured L-band bearer that carries cueing, tactical data link and mid-course updates to the missiles in flight, a Ka-band link to the SAR shell, a Ku-band downlink to ground and a positioning, navigation and timing (PNT) payload.
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: 101 small satellites at ~560 km in three tiers (two mid-inclination revisit tiers and a dawn-dusk search plane) carrying synthetic-aperture radar, giving all-weather detection and revisit across the approaches. Purchased from an allied supplier, Australian-owned and Australian-operated.
JSTN core	The Australian-controlled software and mission-system element of the JSTN that correlates JSTN sensor data into targeting-grade maritime tracks, manages custody and LRPF engagements, provides the weapon datalink and cryptographic services, and exchanges tracks and cues with the wider Defence targeting and command-and-control environment (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 describes the Joint Architecture and illustrates its use across five scenarios, from peacetime to conflict. The scenarios are illustrative rather than predictive and remain 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. First, 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. Second, long-range precision fires in enough depth to make a projection attempt costly across the whole of the approaches, not at a single point. Together, they define the proposed capability.

The architecture is intended to meet this objective with a high degree of Australian self-reliance. 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 is intended to broaden the Government's choices and preserve decision flexibility. It also complements the integrated force and Australia's alliances as an additional near-term layer of denial. It is not proposed as a replacement for existing programmes.

The sensing 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 shared by Australia would contribute directly to that objective while also supporting surveillance of 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. The architecture contributes to four of them:

- **More lethal maritime capabilities.** The primary maritime role is to hold forces at risk across the northern approaches through a land-based magazine that extends anti-ship reach and complements the fleet.
- **Long-range strike.** The Program's direction to "further expand long-range strike capabilities across the integrated, focused force" is addressed by precision fires reaching 2,400 km.
- **Resilient multi-orbit space.** The Program's commitment to secure, multi-orbit space capability for communications is consistent with the space-based custody layer proposed here.

- Autonomous and collaborative fires. The missiles are autonomous and networked: fires can be cued by authorised sensors through the wider targeting environment.

1.3 Capability overview

The architecture delivers the following capabilities:

Persistent surveillance, reconnaissance and targeting (JSTN)

- Persistent space-based RF/ELINT detection, geolocation and classification of emitters (including surface and airborne radars and VHF/UHF communications, across roughly 150 MHz–15 GHz) across the approaches from 50°S to the equator and north.
- **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 4–6 hours.
- On designation, **custody-grade SAR revisit of ≤15 minutes** in the equatorial approaches (+/-20°), tightening to minutes on cue.
- **Spotlight and dwell SAR imaging down to 16 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 operates with multiple sensor-to-shooter paths. A maritime track held by JSTN can be combined with cues and tracks from JORN, ships, aircraft and other Defence sensors through the wider targeting environment, allowing the best-placed shooter to engage. This would support routine Army–RAAF–Navy collaborative engagements, including synchronised time-on-target fires from land-based launchers. With the magazine already forward and cued from space, fires stay responsive around the clock, and redundant sensor-to-shooter paths reduce the risk that a single counterstrike severs the engagement chain.
- **Manufacture at rate.** An Australian government-owned line can build and reconstitute the full 3,600-round magazine within a year. The deterrence argument therefore depends on production rate as well as the standing stockpile.
- **Low manning.** The batteries are containerised, statically emplaced and remotely commanded.

Joint Architecture for Forward Denial — Coverage


Joint sovereign targeting network (JSTN): SAR imaging and ELINT emitter fixes hold custody of maritime targets; ELINT/comms satellites cue massed long-range precision fires (LRPF)



Figure 1.1 The Joint Architecture: coverage. Persistent JSTN surveillance and custody (SAR revisit by latitude, 90th percentile) 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, ~5 m in a calm sea
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	~25 t against an ISO limit of 30.5 t

WHAT A BATTERY DOES



1 battery

Small surface action group

cruiser, frigate, oiler
126 rounds to destroy



2 batteries

Large surface action group

cruiser, destroyer, 2 frigates, oiler
230 rounds to destroy



3 batteries

Carrier group

carrier, cruiser, destroyer, 3 frigates, oiler
339 rounds to mission-kill

THE FORCE

25 batteries · 400 containers · 3,600 rounds

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

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

Figure 1.2 The LRPF battery and its round. Each battery's sixteen armed containers are deployed among 32 to 64 decoy containers (Section 4.4). Illustrative planning estimates; derivations in Sections 3 to 5.

1.4 Scenario one — Task Group 107 custody

In February 2025, 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, but maintaining custody required an improvised national surveillance effort.

With the network in place. 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 at intervals of minutes to tens of minutes by the SAR shell, emitters identified and fixed by the ELINT/comms shell each time a radar or datalink transmits. The resulting JSTN maritime track is published into the wider joint picture alongside JORN, Wedgetail and other Defence inputs. The exercise broadcast is verified by observation rather than relayed by an airliner. No comparable surveillance surge is required: ships and aircraft remain on their programmed tasks because custody is provided as a standing network function. 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. Persistent, visible custody would reduce the advantages of operating at range and demonstrate the surveillance capability without requiring a dedicated shadowing 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, using its presence as visible pressure in a dispute. Coercion requires the group to be seen, to loiter and to stay, and each of those requirements is a vulnerability against a standing architecture. Government would have a sequence of reversible responses, with the early steps intended to impose costs or signal capability without using force against the group.

The response ladder has four rungs. 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 provides a visible signal of attention, although at that distance its immediate coercive effect is limited. A single aircraft or surface combatant would also be operating against a group with its own layered air defence. 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; the movement is therefore observable to the group. Fourth, the scenario considers a single round, fired on an announced track that crosses ahead of the group and terminates in empty ocean.

Purpose of the demonstration. The shot is intended as an unambiguous demonstration. No target is struck and nobody is hurt. It provides observable evidence of four capabilities that a statement alone would not establish. The round's arrival where the group was rather than where it had been establishes that the custody is fire-control grade. Its flight also demonstrates that the reach from Australian soil is real and that the missile works, because it flew the profile and terminated where it was told. Using one round for a demonstration also signals that the magazine is not being treated as too scarce for demonstration. The group can verify these observations with its own sensors and assess their operational implications. The intended

outcome is withdrawal without further escalation, with capability demonstrated through observable performance rather than declaration alone.

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/comms 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/comms 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/comms 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. Section 4 provides the underlying model.

1.7 Scenario four — joint maritime strike with hypersonics

The additional capability. The scenario uses the same architecture but adds HACM, the air-launched hypersonic weapon being introduced through AUKUS Pillar 2 for the Royal Australian Air Force. No official range is published; open sources put it between five hundred and a thousand kilometres.

The network enables the aircraft to engage at range. At HACM's reach the launch aircraft cannot find its own target: its radar detects a warship well short of that distance, and switching it on would warn the group long before the radar found anything. 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 radiating. Nor does the RAAF need to send manned

platforms forward to find the group, inside the reach of its air defences. In this configuration, one network holds the target while several shooters engage it. The network reconciles their different launch points and times of flight to produce a simultaneous time on target, so that hypersonic weapons from the air and long-range fires from the land arrive together.

The engagement. The target is a larger group, or a carrier group with its escorts. The hypersonic weapons target 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 must respond simultaneously to a small number of very high-speed tracks and a larger sea-skimming saturation raid. Combining RAAF hypersonic weapons for air-defence suppression with LRPF could reduce the number of higher-cost munitions required.

Integration of new capabilities. As additional capabilities enter service, the network can combine them with other platforms to enable new forms of joint warfighting. When SSNs enter service, the JSTN can provide timely updates on adversary movements in the open ocean.

1.8 Scenario five — joint land attack with SEAD

The build-up under surveillance. An adversary establishes an expeditionary lodgement on an archipelagic feature within the approaches, some 2,000 km from the Australian mainland: an airfield, a port, dispersed surface-to-air missile (SAM) launchers and the radars that serve them. The site is designated, so the SAR satellites image it in spotlight mode (sub-metre, down to about 16 cm) every few hours through cloud and darkness. Successive images would show changes such as new revetments, movement of transporter-erector vehicles, deployment of surveillance radar and the arrival of aircraft, with individual objects resolvable in the higher-resolution modes. SAR imagery is combined with allied or commercial EO imagery to provide additional assessment. Meanwhile the ELINT/comms 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 survey inputs. 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 provides the required aimpoint precision; 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 seeker types are complementary: 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. Destroying launchers in their revetments and aircraft on the ground would reduce the threats that the follow-on air package must address. 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 RAAF package is substantial. 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. The initial SEAD phase is therefore conducted without exposing aircrew and aircraft to the target-area air defences.

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 recover, crews rest and the tankers prepare for a later sortie.

The scenario illustrates economy of force. Across the whole mission (rounds fired, weapons dropped and hours flown), the joint operation is modelled at roughly half the cost of servicing the same target set from the air alone, because suppression converts expensive stand-off weapons into cheaper direct-attack weapons.

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

The JSTN is a persistent, theatre-wide surveillance, reconnaissance and targeting network, optimised for maritime operations across Australia's approaches and supporting the architecture's shooters. This section defines the network's requirements, presents a design concept, and outlines a fielding sequence. It is a design concept, not an engineering specification. Its purpose is to test whether the required network could be assembled on a relevant timeline from existing technologies. References in this section to AIR 6500 and the wider Defence Targeting Enterprise describe a proposed integration model based on public information; they do not purport to describe non-public Defence architectures, interfaces or program designs.

2.1 Purpose and the joint imperative

The JSTN provides persistent, theatre-wide surveillance, reconnaissance and targeting across Australia's approaches. Its primary operational task is to find, classify, track and maintain targeting-grade custody of surface forces, while also imaging fixed land sites and providing RF/ELINT data on emitting surface and airborne systems. It must do this under Australian authority and resiliently enough to survive an attempt to blind it.

Its joint character follows from how its products are used, not from an attempt to reproduce every domain picture. A denial engagement interleaves domains: a surface action group may operate with airborne early warning, strike aircraft and submarines, while Defence must maintain the wider air, maritime and undersea pictures. Within that joint environment, JSTN provides persistent theatre-wide surveillance, reconnaissance and targeting, optimised for maritime custody and fires. Its maritime tracks and RF/ELINT reports are intended to feed the wider defence targeting enterprise and battle-management systems; relevant tracks and cues from those systems can in turn be ingested by JSTN for engagement context and fires. This makes the JSTN a joint-force contributor without making it the owner of the ADF's complete joint picture.

The scope boundary is deliberate. The JSTN's primary operational focus is the maritime surveillance, reconnaissance and targeting picture required for custody and fires, together with imaging of fixed land sites in support of the fires (Section 2.5). Its ELINT/comms shell also detects, geolocates and classifies airborne emitters, providing a valuable air-data feed, but that does not constitute a complete air picture. Air surveillance and battle management remain functions of Defence's air-battle management systems and sensors. Undersea surveillance and custody belong to the undersea force and its own sensors; JSTN can ingest their tracks when relevant but 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
Persistent theatre-wide surveillance	Persistent maritime surveillance across the five sectors (western, Top End, north-east/Coral Sea, south-east, Indian Ocean island arc), with passive RF/ELINT coverage of emitting surface and airborne systems across the same operating envelope.
Maritime tracks and air ELINT contribution	Continuous surface track production and classification for maritime targets, plus detection, geolocation and classification of airborne emitters as a feed to the wider air picture. The JSTN does not itself produce the complete air picture.
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.
Fusion and C2	Multi-sensor maritime fusion, cryptographic C2, disciplined find-fix-track-target-engage-assess management across dispersed shooters and multiple axes, and interfaces to the wider Defence targeting enterprise.
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 JSTN 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 combines its own space-based surveillance, reconnaissance and targeting functions with interfaces to existing Defence sensors, targeting systems and battle-management systems.

Over-the-horizon radar (external wide-area cueing). JORN provides persistent, very-wide-area north-facing detection of surface and air movement beyond the line of sight [4]. JSTN can ingest relevant JORN maritime tracks and cues, but JORN remains an existing Defence sensor rather than part of the JSTN.

The JSTN space surveillance and reconnaissance layer (SAR shell and ELINT/comms shell). Australian-owned satellites in proliferated low-earth orbit provide all-weather wide-area maritime detection and passive emitter geolocation and classification. The ELINT/comms shell

can detect both surface and airborne emitters, so it contributes useful air-domain data even though the JSTN does not generate the air picture. This layer is developed in full in Section 2.4.

Airborne early warning and AIR 6500 (external air picture and battle management). The E-7A and its successors provide high-fidelity air surveillance and battle management [1]. AIR 6500's Joint Air Battle Management System is publicly described as the command-and-control architecture at the core of Defence's Integrated Air and Missile Defence capability and as connecting information and capabilities from across the ADF [18]. In this concept, JSTN contributes maritime tracks and RF/ELINT reports into that wider environment and can receive relevant air tracks and cues from it. The precise program interfaces are not publicly described and are not assumed here.

The JSTN core (maritime fusion, engagement management and the weapon datalink). The core fuses JSTN's maritime surveillance and reconnaissance data with relevant external cues into targeting-quality maritime tracks, manages custody and LRPF engagements, exchanges track and cue data with the wider defence targeting environment, and carries the mid-course updates. It therefore requires the highest degree of Australian control (Section 2.6).

2.4 The space layer: two shells, two procurements

The concept assumes that the satellites are procured while the JSTN core is developed under Australian control. Both shells are bought from mature allied product lines and operated by Australia. Comparable national systems have been delivered from contract to operational handover within twelve months [10]. A first-of-type Australian spaceborne SAR programme would involve a substantially longer development cycle and would compete for funding and specialist workforce with the fires programme. 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. Ownership removes reliance on a foreign operator for day-to-day tasking and access. Australia builds the JSTN core (Section 2.6); the replenishment cycle every five to seven years is the main opportunity for increasing Australian content once an operational baseline has been established.

Wide-area sensing and continuous connectivity impose different physical requirements, so they are separated into two shells at different altitudes and inclinations and procured separately. Both are sized to the same operating envelope: the approaches from 50 degrees south to the equator, with most operations expected in the priority band from the equator to about 20 degrees south.

The SAR shell

A synthetic-aperture radar (SAR) constellation of 101 small satellites in the ~200 kg class at ~560 km in three tiers (two revisit tiers and a dawn-dusk search plane), bought from an allied supplier and operated by Australia. The specification is based on the current commercial generation: 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 revisit figures in this section assume a fourth-generation (Gen4) class satellite, whose high-resolution imaging strip extends to about 400 km either side of its ground track, against about 150 km for the previous generation, on which the figures would not hold [5].

The shell uses the same SAR for two distinct tasks [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. The search is done by the dawn-dusk plane: at a planning figure of 15 per cent SAR on-time per orbit, its 24 satellites sweep the entire 69 million square kilometre surveillance area about every six hours on their own, or about every four hours crediting the gain in usable imaging time the supplier attributes to Gen4. 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: the two revisit tiers bring any designated point within a steerable look with a mean wait of 5 to 10 minutes across the priority band, 20 to 30 minutes between 30 and 40 degrees south, and about 6 minutes at 50 degrees south. A 30-knot warship needs about 108 minutes to cross a single wide-mode frame, and the longest gap between looks anywhere in the envelope is about 98 minutes, so a cued track is held between revisits.

The full-area sweep is the limiting case. The ELINT/comms shell's passive radio-frequency payload watches continuously and cues the SAR onto any emitting contact. That cue travels through orbit rather than through the ground. Each SAR satellite carries a Ka-band inter-satellite terminal to the ELINT/comms shell above it, so cue-to-look latency is minutes rather than the tens of minutes a ground tasking uplink imposes, and a SAR detection reaches the ground through the upper shell's mesh within seconds rather than waiting for the satellite to pass over a ground station. The terminal is the one non-catalogue item on the SAR satellite. A fully non-emitting contact would instead depend on the hours-scale wide-area sweep. On these assumptions, emitting contacts may be detected within minutes, fully non-emitting contacts within the wide-area sweep period, and designated contacts revisited on a minutes-scale cadence.

Orbit design remains an Australian choice even for purchased spacecraft. The shell flies as three tiers. A high tier of 33 satellites at 53° in eleven planes holds the southern approaches, and a low tier of 44 satellites at 25° in eleven planes serves the equatorial priority band. The 24-satellite search plane flies a single dawn-dusk sun-synchronous orbit, where near-continuous sunlight increases the available SAR duty cycle: a dawn-dusk RADARSAT satellite images 15 minutes per orbit where a satellite of the same mass on a conventional orbit images two [16]. Plane count matters as much as inclination. With an odd number of evenly spaced planes the northbound and southbound ground tracks interleave rather than fall on one another, so 33 satellites in eleven planes give better revisit at every latitude than 36 in six. 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 ($\pm 30\%$), not guarantees.

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

The ELINT/comms shell

A larger constellation of 180 cross-linked satellites at ~780 km, in a single Walker shell of twelve planes at 45° inclination, sized for continuous coverage rather than revisit. Continuous coverage requires substantially more satellites than SAR revisit, supporting the use of a separate shell. The size is set by the datalink. A missile in flight or a small ground terminal sees the shell through a 15° elevation mask, and at that mask 180 satellites keep at least one in view at all times across the envelope, with about three in view in the priority band. At the 5° mask a

passive receiver tolerates, four or more are in view continuously, which supports time-difference geolocation of an emitter and a standalone positioning and timing service. The count also allows the loss of a whole orbital plane without a gap in the datalink, and the 45° inclination concentrates satellites in the equatorial band while still holding the southern edge of the envelope at 50 degrees south.

Iridium provides a relevant reference architecture: 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 at nearly three times Iridium's count and in a mid-inclination rather than near-polar orbit, is one credible near-term 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, not a settled selection, and a competitive acquisition would test whether it can meet the schedule.

Every satellite in the shell carries the same payload complement: L-band radios for the Australian mid-course targeting datalink to the missile, which is the critical function, and for a Link 16 tactical-data-link payload to crewed platforms; the passive radio-frequency and electronic-intelligence antennas and receivers; a Ka-band inter-satellite terminal to the SAR shell below; laser inter-satellite links forming the mesh within the shell; a Ku-band downlink to the ground stations; and a positioning, navigation and timing (PNT) payload, which at this constellation size gives receivers built for it a standalone fix when satellite navigation is denied. Two of those, the electronic-intelligence payload and the targeting waveform with its cryptography, are the strongest candidates for Australian design authority rather than purchase because they are the elements most likely to attract export restriction and the functions Australia would need to modify without relying on external approval (Section 5).

The payload split follows from their physical and electromagnetic requirements. 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; co-locating a sensitive passive receiver with a high-power broadband emitter would compromise both functions. So the radar anchors its own shell, and the radio-frequency payload lives on the ELINT/comms 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 well suited to commercial sourcing.

Two satellite shells, two suppliers, one ground segment

The two satellite shells are bought from mature commercial providers.

The main area of commonality is the ground segment: a common tasking interface, the JSTN core, one cryptographic regime, and one mission training pipeline even where the platform

pipelines differ. The long-lived Australian investment is therefore the common interface and the JSTN core rather than any one satellite generation (Section 2.6).

This approach carries additional costs: two spares pools, two supplier ground systems to integrate, two sustainment contracts, two configuration baselines. This creates a continuing integration and sustainment overhead, but avoids the schedule and first-of-type risks of developing both satellite classes domestically, and both product lines are sustained commercially whether or not Australia buys them.

Frequencies and interoperability

The shells use four 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 inter-satellite link between the shells. Ku-band (~12–18 GHz) carries the ELINT/comms shell's downlink of reports and tracks to ground, and the SAR shell downlinks its imagery on the supplier's own high-rate band. Optical links carry the on-orbit mesh within the ELINT/comms shell.

Reaching munitions and crewed aircraft presents two different integration problems. The munition is the simpler case, because Australia designs it: a purpose-built receiver at the shell's own L-band frequencies, with no legacy constraint to satisfy. Crewed aircraft are more difficult, 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.

A lower-risk approach is to leave the cockpit unchanged and carry a tactical-data-link payload on the ELINT/comms 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 provides practical interoperability: the shell would use the coalition's existing tactical-data-link format and frequency band, allowing a JSTN maritime track, ELINT report or cue to enter the existing joint picture without a new aircraft terminal.

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

Mode	Resolution	Scene	What it is for
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 these SAR modes operate through cloud and at night, unlike optical imaging.

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 supports aimpoint selection within a site rather than simple coordinate targeting.

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. The same comparison supports three separate jobs (indications and warning, aimpoint currency, and post-strike assessment) without requiring a separate sensor system beyond the tasking and orbital phasing already assumed.

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. Keeping these functions distinct also clarifies a potential 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 could support a national research programme involving the Defence Science and Technology Group and universities, with dual-use applications in civilian earth observation. The proposed architecture could incorporate such a capability later, but does not depend on it.

2.6 The JSTN Core: maritime fusion, engagement management and the weapon datalink

The JSTN core is the part of the network requiring the highest degree of Australian control and cannot be treated as an off-the-shelf supplier product.

It comprises multi-sensor fusion of the SAR and ELINT/comms shells with relevant external maritime cues; maritime track management and custody, including cross-phenomenology integrity checks; engagement management across dispersed shooters on multiple axes; interfaces that publish maritime tracks and ELINT reports to the wider Defence targeting enterprise and ingest relevant cues and tracks; 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 reduces the integration and control risks associated with purchasing the satellites from external suppliers.

The JSTN core is primarily a software and mission-systems integration programme. It draws on capabilities already present in Australian industry and retains value across successive satellite generations: the constellations will be replaced twice over the network's life, while the core persists across successive satellite generations and can be incrementally improved.

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. The larger benefit is likely to be in sustainment and continuous development rather than the initial build. 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. That same property underpins the resilience posture in Section 2.7: a continuously maintained baseline reduces the exposure associated with static, infrequently updated software and supports rolling updates.

AI assistance changes the cost structure but does not remove the assurance burden. 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. Faster code production does not reduce that burden; programme planning should assume the same verification and certification effort. The main benefit is greater development throughput and a sustained ability to update the system. That sustained update capacity is important to the architecture and difficult to achieve through conventional acquisition.

2.7 Resilience and reconstitution

Because the JSTN is central to the architecture, it is an obvious target for disruption. The design therefore emphasises resilience and reconstitution.

Hard to blind. The fixed nodes are dispersed, hardened, redundant and designed for rapid reconstitution, so node loss degrades fidelity without severing function, and the multi-layer design means no single sensor type is a single point of failure. The bearer design is intended to degrade gracefully: 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, not megabytes, so the network loses data rate without going dark. Resilience is provided primarily by constellation size, redundancy and rapid replacement, reducing the need to make each individual spacecraft highly survivable.

Escalation implications of physical 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. Physical attacks on Australian or multinational nodes would also carry escalation consequences, which form part of the resilience context. Commercial feeds and bearers are a resilience complement, not the assured layer; the Australian nodes remain the backbone.

Cyber attack presents a different resilience problem. Unlike physical attack, cyber activity can remain deniable and below the threshold of armed conflict; it is also continuous, and the coercion campaign at Section 1 opens with it. The network is therefore designed for persistent cyber attack, with 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 also 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 the threat rather than changing its nature: machine-speed vulnerability discovery will probe the attack surface faster than any static defence can be certified. The response is continuous iteration of the JSTN core (Section 2.6) and, through it, the fleet, with the same machine-speed methods used for defensive monitoring.

2.8 Building on existing investment

The JSTN is achievable because it adds persistent, theatre-wide surveillance, reconnaissance and targeting, optimised for maritime operations, to capabilities the ADF already fields rather than recreating the wider joint picture. The main new investments are the JSTN core and the two space shells; existing Defence sensors and battle-management systems remain external systems with which JSTN must exchange data.

Table 2.3 *Existing Australian investments and their role in the integrated architecture.*

Existing investment	Role in the integrated architecture
Jindalee Operational Radar Network	External wide-area surveillance and cueing; relevant maritime tracks can feed JSTN
E-7A airborne early warning and control	External high-fidelity air picture and battle management; relevant air tracks and cues can be ingested by JSTN
Integrated air and missile defence C2	Wider joint air-defence command-and-control environment; AIR 6500 provides the JABMS integration path described in this concept [18]
Defence and commercial ISR and satellites	External space-based sensing augmenting the Australian-owned maritime-custody layer
New: the JSTN core	Maritime fusion, custody, LRPF engagement support and weapon datalink; interfaces to the wider Defence targeting enterprise

Institutional approach. The JSTN should be developed as a joint-force surveillance, reconnaissance and targeting component of the wider defence targeting enterprise, optimised for maritime operations, with LRPF integration from the outset and defined interfaces to existing sensors and battle-management systems. AIR 6500's JABMS provides a public example of the wider joint C2 environment into which JSTN maritime and ELINT products could feed [18]; the exact organisational and technical boundary is not publicly specified. 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]. This would extend an investment category the Government has already established.

2.9 Sizing and cost: the trade space

The space-layer cost is a trade space across the two shells, their satellite counts and the coverage each buys. Published procurements provide the cost anchor.

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 provides a useful comparison between system ownership and a data-and-operations service [10].

Those transactions suggest two cost implications. 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 national segment (ground stations, tasking, the JSTN 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. The ELINT/comms satellite is a heavier spacecraft, about 520 kg against about 200 kg, with a larger payload complement; it is priced from communications-constellation market anchors plus the payload difference, at about A\$60 million delivered and launched. Second, small-fleet system prices cannot be multiplied linearly to a hundred satellites, because the fixed segment is amortised across the larger fleet and production-scale pricing changes the marginal unit cost (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, 101 satellites (baseline)	National; mean wait to next look ~5–10 min in the priority band, ~20–30 min at 30–40°S; full-area detection sweep ~4–6 h	~A\$4.0–4.5 bn	All-weather search, custody and identification of the approaches and the whole continent
ELINT / communications, 180 satellites	Continuous, 50°S to the equator	~A\$11–12.5 bn	The Australian fire-control link, emitter geolocation and PNT, with optical cross-links within the shell
Combined space layer	—	~A\$15.5–17.5 bn	Both shells and the fixed national segment

These are rough-cut planning estimates. The larger shells of Section 2.4 have been priced by scaling the per-satellite, launch and non-recurring lines to the new counts, and a full probabilistic re-estimate is still to be run. Version 1.1 of this volume printed A\$3.2–3.8 billion for a space layer of 110 to 126 satellites. The increase follows from the larger constellations and the heavier ELINT/comms satellite, and it is carried because the two constellations are the enabler of the whole architecture. For scale, the space layer at this sizing is about 4 per cent of the A\$425 billion Integrated Investment Program decade [12]. JP 9102, the geostationary SATCOM system Defence cancelled in November 2024 in favour of a resilient multi-orbit

approach broadly consistent with the posture proposed here, remains the relevant policy precedent (quoted at Section 5.2) [11].

2.10 Development sequence, and the interim

The JSTN is fielded incrementally, so each step adds usable targeting capability before the final configuration is complete.

1. Establish the JSTN core and its interfaces. Build the core against commercial SAR/RF feeds and available Defence maritime cues, including JORN, while defining and testing bidirectional interfaces to the wider Defence targeting enterprise and the AIR 6500/JABMS environment. The core should be able to publish maritime tracks and ELINT reports and receive relevant external tracks and cues without making AIR 6500 a prerequisite for an LRPF engagement. This work starts first because it takes longest to mature and because the core must remain under Australian control.

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/comms shell. Scale the SAR constellation toward its full size while the ELINT/comms shell, the electronic-intelligence payload and the Australian targeting datalink mature behind it. The JSTN core is expected to be the longest-lead element and is sequenced accordingly.

4. Replenish, and grow the Australian content. The first replenishment cycle provides an opportunity to increase Australian content.

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 serves two purposes. It provides SAR custody over the approaches years before the owned satellites fly and gives the JSTN core a live operational feed to be built and tested against, in the products and formats the owned shell will later deliver. The core therefore matures on operational data before the owned constellation is delivered and reduces integration work at handover: using the eventual constellation supplier keeps the core on the same interface. The subscription provides access rather than sovereign 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.

Interim capability and 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 interim approach should remain diversified: 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/comms shell; commercial as diversified backup bearer

The providers are included only to demonstrate that each layer is currently procurable from allied sources; this does not endorse any one. Capabilities, ownership and availability change, and several are subject to export control or supplier-government priority in a crisis. These constraints support their use as an interim and backup layer, with the owned shells providing the assured layer.

References

All references were accessed in July–September 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 external sensing and battle-management source whose relevant tracks and cues can be exchanged with 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/comms 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, cited as an external wide-area cueing and tracking source relevant to JSTN maritime fusion.

[5] Small-satellite SAR design and imaging references: ICEYE X-band constellation (Finland), 85 kg (Gen2) to ~200 kg (Gen4); Gen4 widens the high-resolution access strip from about 150 km to 400 km either side of track and offers resolution to about 16 cm (ICEYE releases, March and September 2025); 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); Synspecive (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.

[18] Australian Department of Defence, 'Joint Air Battle Management System', Project AIR 6500 Phase 1; Australian Government, Defence Ministers, '\$2.4 billion investment in air and missile defence', 9 September 2026. Public descriptions identify JABMS as the command-and-control architecture at the core of Defence's Integrated Air and Missile Defence capability and as connecting information and capabilities from across the ADF. Cited only to establish the public role of AIR 6500/JABMS; the specific JSTN interfaces proposed in this volume are architectural assumptions, not public Defence design information.

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

The JSTN provides targeting; the Long-Range Precision Fires provide the strike component. This section defines the fires requirements, missile design, launch arrangement and flight profile. The effectiveness model is at Section 4; the industrial path is at Section 5.

3.1 Requirements

Five requirements set the design and, taken together, make an off-the-shelf purchase unsuitable.

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 that emplaces itself on its own legs; 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.

No system currently in or entering service meets the first two requirements 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, a unit cost that is difficult to reconcile with the massed-salvo concept modelled here. Section 5 sets the two production models against each other and prices the gap.

The last two requirements shape the industrial model. 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, cannot be guaranteed through procurement alone. For that reason, the concept assumes an Australian-controlled design and production line. This does not preclude allied integration: 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 no detected track can safely be ignored. The architecture assumes no dedicated decoy rounds: at the modelled US\$300,000 unit cost, each additional round can both consume defensive capacity and deliver an effect. The effect is modelled at Section 4.

The same round serves land targets. For targets with identifiable critical points, precision can reduce the volume of fire required, 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, aimed 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 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 2,400 km specification includes reserve against adverse conditions: 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.

The selected range is a policy choice. 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. The selected range affects which targets the force can credibly hold at risk and should therefore be set by Government against declared deterrence policy rather than by the maximum physical range of the airframe.

Engine and airframe are chosen for cost. The design uses a turbojet rather than a turbofan because it has fewer parts and stages and is cheaper at scale, at the price of a larger fuel fraction, which is acceptable for a missile produced in thousands. Small turbojet manufacture is now available from a wider commercial and defence industrial base than in earlier generations. 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 manufacturing method is selected for lowest cost at production volume while meeting low-observable and structural requirements; it does not require the manufacturing approach used for 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 area fragmentation is more useful than a point-impact effect. 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 directed at spaces whose loss would most affect combat capability. The first aimpoints are 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]. The command spaces come next, followed by 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. Navigation uses satellite and inertial systems, 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 key development 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 that stands on its own telescoping legs and travels on a flatbed trailer. The dimensions are taken from published ISO data [8], because launch density determines battery size and ultimately force structure.

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

A three-across by three-high arrangement gives **nine rounds per container** in a 3 × 3 array. The high-cube is required: 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 flatbed trailers rated for the 40-foot high-cube at that gross mass. The container raises and lowers itself on its own legs (Section 4.4), so no side-loader or crane is needed, and the legs fit within the margin of about five tonnes to the ISO maximum.

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. The concept assumes static, dispersed emplacement rather than shoot-and-scoot operation. 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. Disabling a battery would therefore require a substantial allocation of long-range fires or the use of platforms operating well forward.

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 may be more suitable for forward basing than high-value platforms because the loss of an individual container represents a smaller fraction of overall force capability. Manus is sovereign PNG territory, so the option is contingent on host-nation agreement.

3.4 The attack profile

The raid depends on networked targeting through to the terminal engagement. Before interception begins, effectiveness depends on the raid maintaining custody and updating its plan faster than the group can reposition or extend its warning screen. The flight profile is designed around that networked targeting process.

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, rather than an optional tactical refinement.

A horizon-hugging descent. 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. This preserves cruise efficiency while allowing the descent to contribute to terminal speed.

Airborne radar produces a different horizon and requires a modified profile. 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.

Early warning does not itself provide interception capacity; the air wing is therefore treated as a finite 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 sharply reduces the utility of 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 main departure from the fielded pattern is the passive seeker, which does not radiate during the terminal approach.

The preferred terminal approach is against the ship's 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 and should be re-verified before later use. 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, then uses the results to derive the rounds required, battery and force structure, and 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 case for the missile.

4.1 The effectiveness model

The model begins with cell counts. A defended surface group's air-defence capacity can be bounded using countable quantities: 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 every round in the salvo is a genuine weapon, the defender must treat each detected track as a potential threat.

These constraints allow the magazine-exhaustion threshold to be estimated without relying on a single assumed kill probability. The model is built this way because cell counts are checkable from public sources, while a full chain of per-layer kill probabilities is much less certain. Higher kill probabilities can extend magazine endurance, but do not increase the number of available interceptors.

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 and the simultaneity ceiling a single coordinated salvo imposes, even though the latter 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. These assumptions bias the estimate toward the defender. The aim is to avoid relying on favourable assumptions about those uncertainties.

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. The model deliberately takes the defender-favourable reading and prices the group 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. This is a modelling margin: a machine-vision seeker aimed at critical spaces wastes fewer hits than this, so the actual requirement could be lower.

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 follows the configuration observed in the Task Group 107 deployment. The carrier-group 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.

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

Type 055 area cells assumed	Small SAG rounds	Large SAG rounds	Carrier group rounds	Battery ladder
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 operationally unrealistic 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. The force-structure result is therefore not sensitive to the precise loadout assumption within the realistic range modelled.

The attacker-side parameters have a smaller effect on the result 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. The reason is structural: the early rounds are spent emptying the magazine whatever their accuracy, so accuracy enters only at the margin. The result is therefore not highly sensitive to optimistic seeker performance.

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 modelled separately from a group. 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, whose limit is the number of engagements it can run at once in the short window before impact, not the number of missiles it holds. 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 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. An undersized raid does not produce a proportionately smaller effect: if it stops short of exhaustion, the defended group remains afloat and alerted. Reserve batteries provide margin if the assumptions prove optimistic, which is also why the force structure is expressed in whole batteries rather than individual rounds.

4.3 Battery and force structure

The force structure is derived from the container, with dispersal at each level intended to prevent a single counterstrike from removing a significant fraction of the force.

Why sixteen containers. The battery is sized so that the three threat cases fall on clean multiples, and remain close to those multiples across the sensitivity range considered in Section 4.1.3:

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. The battery size was selected so that these representative cases correspond to whole-battery increments.

Force-level engagement capacity. Twenty-five batteries is twenty-five small-group engagements, or twelve large-group engagements, or eight carrier-group engagements, before any reload. Sustained operations depend on the production line behind the standing magazine.

Force size and distribution. 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 an existing security perimeter, tenure that does not lapse, and Defence's own works approvals in place of eight state and territory planning systems. It also provides the separation distances required for explosive ordnance: 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 critical to overall coverage: 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 artillery is the Army's established organisation for land-based fires. 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. Manning is in the tens per 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.

Security is based on remote monitoring and a central reaction element. A battery is sixteen armed containers and between thirty-two and sixty-four decoys, two to four for each armed container, 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. The decoy ratio is set site by site: a higher ratio adds positions an attacker must engage, and adds site works and patrol effort in proportion.

Site works are kept to a minimum. Each container stands on four telescoping legs at its corner castings, with broad foot plates that spread its load onto levelled natural ground, so no position needs a prepared hardstand. The legs are electrically driven, like the elevation gear, and raise the container clear of its trailer so that the truck drives out from under it; the same legs lower it back onto a trailer when it moves. The trailers are therefore standard flatbed container trailers rather than side-loaders, and no crane or other lifting equipment is needed at the position. What each position still needs is a cleared and graded spot that drains, lightning protection and earthing, and a fire break. Decoy positions need the same and stand on the same legs, so live and decoy positions look alike on the ground. With nothing built at a position, moving a container within a sector is a transport task rather than a works project.

Ground conditions and access are the remaining variables. At about 25 tonnes gross, each foot carries some six tonnes, which a plate of about a square metre spreads to around 60 kPa, within the bearing capacity of firm natural ground. In reactive clays and black-soil country that soften in the wet, bearing mats under the feet do the work of a pad, and the legs re-level the container as the ground settles. All-weather access is the main site-cost variable and can be considered separately from launch availability: 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 planning allowance is of order A\$10 to 30 million per 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 are likely to drive the schedule more than 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. Keeping ground disturbance to clearing and grading narrows the scope of the environmental and heritage assessments without removing them, and leaves explosives licensing unchanged. Site selection should begin with the development programme rather than with the first delivery.

4.5 Cost-exchange, reload asymmetry, and the line

The model produces a large cost differential. One battery is 144 rounds at US\$300,000, some US\$43 million. Against a small SAG with an estimated replacement value of US\$1.45 billion, the ratio is about 38:1; against the modelled carrier group it is about 74:1. These ratios also exclude the further asymmetry that the fires are sovereign and resuppliable while the ships are not replaceable in any relevant timeframe. On the model's assumptions, the raid remains affordable even with substantial inefficiency, supporting the proposition that the capability could be employed repeatedly.

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

A group that successfully defeats the raid may nevertheless be unable to continue the deployment at the same level of protection. 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 defeated raid could therefore still produce an operational effect by forcing withdrawal for reload.

This creates an asymmetric failure case. The attacker's failure case is a delay and US\$43 million, while the defender's success case is an empty magazine, a long transit, and a station that is now uncovered. Any successful defence still consumes interceptor inventory, although the extent of that effect depends on the defender's actual loadout and engagement doctrine.

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. Each escort added to a group raises the required raid by roughly 45 rounds for a frigate and more for a heavier ship, or some US\$14 to 23 million, against the several hundred million dollars the escort costs to build, crew and sustain. In force-structure terms, that is a cost-imposition ratio of roughly twenty-to-forty-fold against the defender.

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 consistent with the production model now used in United States affordable-mass programmes and with emerging Barracuda local-production targets discussed in Section 5.4; it

is not a quoted price for an existing missile of the same range and payload. 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.

Production rate underpins the inventory. The 3,600-round magazine provides the initial buffer, while a warm production line is intended to rebuild the full inventory in about a year at surge rate. Because capability lives in software, banked rounds take the line’s upgrades rather than ageing in their containers, so stored rounds can receive capability updates during maintenance rather than becoming progressively obsolete.

Vertical integration and unit cost. The affordability argument depends 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; Section 5 sets out the acquisition structure intended to preserve this.

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]. What remains untested is the combination of these factors at scale against a warship, and the force treats it as untested: over-matched sizing, sufficient reserve to absorb modelling error in an initial engagement, and terminal parameters measured rather than assumed.

Each counter increases the required raid size. Each falls into one of three families, and each works by raising the size of the raid required rather than eliminating the underlying magazine and cost asymmetries.

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.

Counter family	Examples	Design answer
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.

Compounding pressures on the defence. 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, continue operating as sensor damage accumulates, 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 effects can compound rather than occurring independently. Combat-management overload is a second exhaustion point that can arrive before the magazine empties, and the model assigns no additional effect to it, keeping the estimate conservative.

What would falsify this section. Two developments would materially change the conclusion. First, a directed-energy defence able to sustain enough concurrent engagements to match the raid arrival rate, despite power, fire-control and marine-atmosphere constraints, would break the rate-limited case. Second, a short-range missile defence with enough fire-control channels and interceptor depth to sustain engagements at the raid arrival rate would weaken the model's defensive bottleneck. Neither capability is evident today; either would warrant revisiting the model and force structure.

References

All references were accessed during preparation of this volume in July 2026. Several of the systems cited are in rapid development and should be re-verified before later use. 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 to assess whether the industrial base required for this architecture is identifiable and how much of it is already being established 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 can be graded by how much of the design, intellectual property and production Australia holds. The architecture applies a different degree of sovereign control to each capability, 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. A key test is whether Australia can resupply the capability in a crisis without requiring external approval; this informs how far each element is taken toward sovereign production.

The 2026 Integrated Investment Program provides the policy framework [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”. Each industrial path proposed in this section aligns with a priority already identified and funded by government. The section applies that existing industrial policy to the proposed architecture.

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

For the JSTN, the industrial task is to add the surveillance, reconnaissance, targeting, communications and Australian-controlled integration capabilities not already in service. JORN, E-7A and the IAMD/AIR 6500 environment are existing Defence capabilities with established sustainment bases; they are treated here as external systems with which JSTN must integrate, not as an in-service spine owned by JSTN. 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]. The architecture specifies this ownership-and-operation model. ICEYE is one candidate supplier; the requirement is for an Australian-owned and operated constellation, with supplier selection remaining open.

The ELINT/comms constellation: purchased, with jointly developed payloads. The ELINT/comms shell carries custom payloads (the ELINT receivers, the assured L-band fire-control bearer and the PNT payload 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 place an operational global L-band constellation, internationally coordinated spectrum and replacement launch within one supplier group: 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 remains a supplier and reconstitution partner, preserving support for replenishment without giving up Australian operational control.

The JSTN core. This Australian-controlled software and mission-system element comprises the basestation interfaces, the maritime data-fusion and custody engine, AI imagery and signals analysis, interfaces to the wider Defence targeting and battle-management environment, and the weapon datalink. Find-fix-fire authority for LRPF must remain 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 JSTN 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], providing an existing sovereign base in SAR processing and analytics.

Recent Defence decisions provide a relevant acquisition precedent for the proposed multi-orbit approach. JP 9102, the geostationary SATCOM program for which Lockheed Martin was selected in 2023 and which was commonly reported at about A\$7 billion, was cancelled in November 2024. Defence stated 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 produced on shared component bases rather than as bespoke prime-vendor developments (Glossary, GOCO; the delivered-price benchmarks are in Section 2.9). The JSTN's proposed acquisition model follows that multi-orbit direction.

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]. Several retained tiers are already in production or construction within that enterprise.

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]. The export component supports economic production at that rate. A solid-rocket-motor manufacturing complex (Mulwala) and explosive-assembly facilities (Orchard Hills) are funded to close the component gaps.

This base provides three specific elements for the sovereign round: 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. It does not, however, supply 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 on the assumptions used here, extending those licensed lines would not by itself achieve a US\$300,000 round.

Section 5.4 therefore examines a separate production path: a sovereign missile and the production organisation needed to achieve its price.

5.4 The sovereign missile: a mature class, produced differently

The missile is the strongest candidate for full sovereign design and production because the airframe is conventional, the intellectual property could be held domestically, and assured resupply is central to the case for an Australian massed-fire capability. 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]. The section addresses four points: 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, North Korea and Pakistan also field or claim weapons in this class. The class has been developed by a range of industrial economies, and Ukraine provides a recent example of rapid wartime development. The distinctive requirement is producing the missile at Table 5.4's price; the main unresolved industrial issues are therefore production method and cost control rather than basic missile physics.

Representative example and candidate collaboration: Destinus Ruta. Destinus provides a current example of this class on a modern production model. Its 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, with a similar emphasis on sovereign supply.

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. Together with the fielded systems at Table 5.2, it shows that the proposed missile belongs to an established class rather than representing a first-of-type technology.

The requirement is price

Comparison of production models. The architecture's unit-cost assumption depends on a different production model rather than a discount to prevailing prices. 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 used throughout because they provide the more conservative comparison.

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

Volume-manufacture model	Unit cost	Scale
Barracuda-500 (Anduril)	~US\$215,000	thousands
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
LRSM 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 a wartime reference point: 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.

Barracuda as a cost and manufacturing anchor. The Barracuda family provides a more specific recent reference point. Taiwan's NCSIST has stated a target below NT\$6.5 million (about US\$216,500 at the time) for locally produced Barracuda-500 rounds under technology transfer [21]. In the United States, Anduril's May 2026 production agreement covers at least 3,000

surface-launched Barracuda-500M systems over three years; Anduril states that the design uses 70 per cent commodity components and can be assembled in about 30 hours with ten common hand tools [22]. Poland's July 2026 agreement to assemble and then produce thousands of Barracuda-500M rounds in a state-owned industrial base, with progressively increasing Polish and European content, provides a directly relevant allied precedent for transferring the low-cost production model [23]. Latvia's September 2026 move toward acquisition adds another allied procurement case, although no quantity or price has been disclosed [24]. These systems are smaller and shorter-ranged than the proposed LRSM and are therefore manufacturing and cost anchors, not direct unit-cost comparators.

The gap between Tables 5.3 and 5.4 approaches an order of magnitude, and the difference is attributable primarily to production method rather than missile physics: much of the guidance, sensing and compute content is now commercially available, and what remains expensive is engine, airframe and militarised qualification. The ~US\$300,000 LRSM figure is therefore a build-to-cost planning target interpolated against this emerging volume-manufacture evidence, not a quoted price for an existing 2,400 km, 250 kg-warhead missile. The remainder of this section sets out the production practices and organisation intended to hold that target.

Production engineering

Because the class is mature, most of the required innovation is in production engineering. Four practices, each demonstrated in the Table 5.4 programmes, contribute materially to the target cost [10][12].

The engine is additively manufactured. The turbojet has been a historical cost and supply constraint in cheap cruise-missile production, and additive manufacture is creating additional options for reducing cost and lead time: 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, with the choice determined by schedule, cost and technology-transfer terms. If a design is purchased, the concept assumes transfer of sufficient rights for domestic production: 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 sovereign-production concept is intended to reduce.

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 are designed into the

airframe from the outset; it is not converted from a hand-built prototype. The design also adopts consolidation practices common in automotive manufacturing: single castings replace fabricated assemblies of dozens of parts, adapted here to missile production. 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 production rate responsive because output scales with demand by adding shifts and cells rather than by recruiting and certifying a workforce. This responsiveness supports the war-stock concept: the line and its demonstrated surge rate are an important part of the capability alongside the inventory 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 updates are not tied to hardware blocks. 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 follows the enterprise's existing trajectory, 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.

Production organisation and cost

The organisational model behind the factory is central to achieving the prices in Table 5.4. 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, with no proprietary element: 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 primarily by control of key functions and assets. 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 production options are assessed against the holdings above and the Table 5.4 cost target. 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 option provides the strongest formal sovereign control, but has the longest start-up 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. Option B offers schedule and development advantages: 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]. Its trade-offs include a manufacturer whose parent is United States-owned, with the associated United States export-control constraints, and a predictable negotiation over the software layer, because that layer may contain commercially sensitive intellectual property. The architecture limits that exposure structurally: most targeting and mission-management functions remain in the JSTN rather than onboard 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
JSTN core (targeting software and 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
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 two options can be assessed against common evidence-based criteria: time to the first block at rate, the surge rate as demonstrated rather than certified, whole-of-life cost including the capacity fee where one is paid, and the five holdings above, which remain constant. Under either option, the contract structure must support volume pricing. A conventional high-cost platform acquisition model would be inconsistent with the unit-cost target, while inadequate Australian control of the five holdings would weaken the self-reliance objective.

The sovereign development items

The main sovereign development items are the same under either production option.

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

Element	Path & candidate partner	Note
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 critical development item, and Australia has relevant domestic capabilities. 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 a relevant RF engineering base: detecting, characterising and homing on a hostile emitter is a radio-frequency sensing-and-processing problem within Australia's existing CEAFAIR and CEAMOUNT radio-frequency engineering expertise.

Recent combat use provides evidence that low-cost machine-vision terminal guidance can operate under jamming: 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 provide external reference points and possible collaborators; combined with the domestic engineering base, they support the feasibility case for an Australian-developed seeker.

5.5 Summary: the base largely exists

Overall, much of the industrial base required by the architecture is already being established under priorities identified in the 2026 IIP; the proposal would need to direct and extend that base rather than create it from first principles. The Defence sensing and battle-management systems with which JSTN interfaces are 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 Australian autonomy firms and SMEs provide a growing domestic base for high-rate uncrewed and guided systems.

The main early implementation priorities are to build the JSTN core, bring the two constellations into Australian ownership and operation, and stand up the sovereign missile on an Australian line under one of the two production options in Section 5.4.

Three qualifications remain. GWEO is a phased programme, and some critics argue that its delivery rate needs to increase against the strategic 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 relevant industrial base is identifiable and, in several areas, already operating or under development.

References

All references were accessed in July–September 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; CEAFAAR 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.

[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.

[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.

[21] Reuters, ‘Taiwan shows off first missile to be jointly manufactured with US arms maker’, 17 September 2025: Taiwan’s National Chung-Shan Institute of Science and Technology stated a target unit cost below NT\$6.5 million (about US\$216,500 at the time) for domestically produced Barracuda-500 missiles following technology transfer. Cited as a local-production cost target for a smaller affordable-mass cruise missile, not as a direct LRSM comparator.

[22] Anduril Industries, ‘Anduril, Department of War Sign Production Agreement for Surface-Launched Barracuda-500M’, 13 May 2026: framework for a minimum 3,000 all-up SLB-500M systems over three years; Anduril states 70 per cent commodity components, the remainder de-risked through competing open-architecture designs, and approximately 30 assembly hours using ten common hand tools. Cited for production architecture, labour content and programme scale; the agreement does not disclose unit price.

[23] Polish Ministry of State Assets, ‘PGZ i Anduril uruchomią w Polsce produkcję pocisków Barracuda-500M odpalanych z wyrzutni naziemnych’, 6 July 2026: agreement for WZL-2 and Anduril to establish assembly and then production of thousands of SLB-500M systems in Poland, with progressively increasing Polish and European component content. Cited as an

allied precedent for transferring a low-cost cruise-missile production model into a state-owned domestic industrial base.

[24] Thomas Newdick, 'Latvia To Order Anduril's Barracuda Cruise Missile For Long-Range Strike Capability', The War Zone, 25 September 2026: Latvia moving toward acquisition of the Barracuda family and the article's discussion of the Polish production precedent; the Latvian announcement did not disclose variant, quantity, delivery schedule or price. Cited as evidence of allied procurement interest and distributed European industrial participation, not as unit-cost evidence.