# The Launch-Range Resilience Decision Gap

*Why the most important question about Cape Canaveral isn't about a wall.*

## We're Asking the Wrong Question

Every conversation about protecting Cape Canaveral from rising seas starts the same way: "*How high does the wall need to be?*"

It's the wrong question. The right one is: **"How much assured access to space do we buy per dollar — and which resilience architecture actually survives the next 50 years of water?"**

A seawall is a single point of failure. The ocean doesn't have to climb it — it only has to outlast it. We don't engineer anything else this way. We don't run a single data center and call it resilient. We don't route all network traffic through one cable. Yet we treat the nation's primary launch range as a fixed asset to be walled off, rather than a launch capability to be distributed.

## The Gap Nobody Is Filling

The Department of Defense already mandates sea-level scenarios for its coastal installations through the Regional Sea Level (DRSL) framework. Installation planners know their exposure. What they don't have is a way to compare hardening, distributing, and relocating a launch range on the only three metrics that matter at budget time: lifecycle cost, mission availability under uncertainty, and resilience across a 50-year planning horizon.

That gap is where billion-dollar decisions get made by whoever argues loudest. The result is a debate about walls when the real conversation should be about **assured-access-per-dollar — a score, not an opinion.**

## What a Defensible Answer Looks Like

A decision engine that scores competing resilience architectures against the same mandated scenarios would let a program officer answer, with evidence:

- Does $400M in distributed launch buy more assured access than $400M in hardened seawall?
- At what point does the cost of delaying a decision exceed the cost of picking any option today?
- Which architecture portfolio still works if the worst-case DRSL scenario materializes by 2060?

These are answerable questions. The data exists. The gap is the model that connects them.

## Why This Is Urgent (and Not Because We Said So)

Construction costs compound — and in the specialized coastal and marine construction market, escalation has run well above the general building index. Every year of delay quietly removes the lower-cost architectures from the menu: an option that could be permitted and built today costs meaningfully more for each year it waits, and a sea-based option with a multi-year development cycle simply can't be in the water when you need it if you start the clock late.

> **The expensive part isn't acting. It's deciding late** — after escalation has compounded and the cheapest options have already dropped off the table. Indecision isn't free. It's just an invoice that arrives later.

## Who We Are and Why We're Saying This

MilkyWayEconomy is an economics advisory firm. We model infrastructure cost and risk — including the public cost estimates for protecting Cape Canaveral that have circulated in the space community. We hold an active federal research relationship in this domain. We don't sell a wall. We don't sell a launch platform. We build the decision models that let someone else build the right thing.

This brief is the public version of a conversation we think the space-launch community needs to have. No proprietary data. No scare numbers. Just the argument — and an open invitation to debate it.

**Read the framework or open the debate with us → official@milkywayeconomy.com**

**MILKYWAYECONOMY**
MilkyWayEconomy.com · official@milkywayeconomy.com