



PROPOSAL

Put storage.ca on the engine we already built

We spent this year building a ranking portal for plumbing.ca. It already runs 19 different lists. storage.ca is still a single WordPress box that can't generate a page for the one thing customers search for — unit sizes.

The ask

Reuse the portal for storage.ca rather than maintain a second, weaker system.

The cost

Close to nothing in infrastructure. The engine is built, paid for, and already multi-tenant.

The catch

One genuinely new piece of work — storage units — and two access questions we need answered.

storage.ca today

A WordPress site on one small server, launched January 2024 and largely untouched since.

What's published

BUSINESS PAGES	932
CITY PAGES	12
UNIT-SIZE PAGES	0
TOTAL URLS	951

Counted from their own sitemap today.

What it runs on

SERVERS	1
SIZE	2 CPU / 2 GB
REGION	Canada
LAST BACKUP	23 Jan 2024

No redundancy. If that one machine fails, the site is down.

Worth pausing on: the last backup was taken the day it launched, twenty months ago. Everything published since then exists in exactly one place.

plumbing.ca isn't a website — it's an engine

The thing we built for plumbing.ca generates ranked, city-by-city lists automatically. It is already running far more than plumbing.

Live on it today **19 LISTS**

- Plumbers, roofers, HVAC, electricians
- Waterproofing, dentists, walk-in clinics
- Event venues, injury lawyers, real estate

What it does per list

- Pulls every business in the chosen cities
- Scores and ranks them on our own criteria
- Publishes a fast, static page per city
- Refreshes on a schedule, no one touching it

Adding a category is a config change, not a project. That's the whole argument for reusing it — we've already paid the build cost once.

We have no pages for how people actually shop

Someone looking for storage searches *"10x10 storage unit Toronto"* — a size, in a city, with intent to book. We publish nothing that matches that.

SEARCH INTENT	PAGES WE HAVE	PAGES THE ENGINE COULD PUBLISH
A named facility	932	932 – kept
Storage in a city	12	12, ranked and refreshed
A unit size in a city	0	12 cities × sizes
Price comparison in a city	0	one per city

The zero in that third row is the opening. Those are the highest-intent searches in this category — the person is ready to book — and today they land on a competitor.

Honest caveat: I don't have search-volume or current traffic figures — those live in Search Console, not AWS. Worth pulling before the 3:30 so the opportunity is sized, not just described.

What we get beyond the pages

New cities become free

Adding Calgary or Halifax is a line of config. On WordPress it's someone hand-building pages for a week.

Rankings stay current

The engine re-scores on a schedule. Our storage pages are a 2024 snapshot that ages every day.

One system, not two

Every improvement to the portal — speed, layout, schema — lands on storage.ca automatically.

Pages get fast

Static files from a CDN instead of WordPress rendering per visit. Speed is a ranking factor.

The strategic point: we're currently paying to maintain two ways of doing the same job, and the weaker one is on the bigger domain.

Three things we're carrying right now

- 1

No backup in twenty months **URGENT**

Two snapshots exist, both from launch day in January 2024. A disk failure loses everything published since. This is a one-command fix and I can do it today — it needs no decisions from anyone.
- 2

We don't control the storage.ca domain **BLOCKER**

The DNS settings we hold for storage.ca are inactive — the live ones are managed somewhere outside our AWS account, as is the CDN in front of it. We can't move the site until we know who holds them. This is the single biggest unknown in the plan.
- 3

Someone may already be building this **CHECK**

There's a separate storage environment in our account that was updated yesterday. If that's the same initiative, these should merge before either goes further.

What's actually new: units

Every category on the engine today ranks *businesses*. Storage is the first one where the customer compares what's *inside* the business — a 5×10 against a 10×20, by price.

What the engine knows now

One record per facility: name, address, rating, reviews. Enough to rank a plumber. A facility with twelve unit types collapses into a single row.

What we'd add

A record per unit type per facility — size, square footage, price, climate control, availability. This is what unlocks the comparison pages and the size-based searches.

This is the one open question with real weight: where does unit and pricing data come from? Facilities supplying a feed, scraping their sites, or manual entry. That choice drives most of the timeline — everything else is assembly.

How I'd sequence it

now

Protect what's live

Turn on backups for the current site. Find out who controls the domain and CDN. No dependencies — starts today.

step 1

Prove the ranking works for storage

Point the existing engine at one city and see whether the list it produces is credible. Cheap, reversible, and tells us early if storage needs different criteria than plumbing.

step 2

Build the units piece

The only genuine engineering. Scope depends entirely on where the data comes from.

step 3

Run both sites side by side

New version live on a test address while WordPress keeps serving customers. Nothing switches until we can compare them directly.

step 4

Switch over

Move the domain once the new pages match on content and SEO. Old site kept available for a while, then retired.

On timing: steps 1, 3 and 4 are days of work each. Step 2 could be a week or a quarter depending on the data-source answer — I'd rather scope it after step 1 than guess now.

What it costs

The infrastructure argument is easy: we're not buying anything. The engine, its load balancer and its storage are already running and already shared across 19 lists.

ITEM	STATUS	ADDED MONTHLY COST
The ranking engine	Already running, shared	none
Page hosting + CDN	Already running, shared	negligible
Unit data storage	New	small
Current WordPress server	Retired after switchover	saving

The real cost is engineering time on the units piece, and that number depends on the data-source decision.

So the case isn't cost savings — the server is cheap. It's reach, resilience, and not maintaining two systems that do the same job.

What I need to move

Q1 Who controls storage.ca's domain and CDN?
Sets the switchover path and most of the schedule. Everything else can proceed while we find out.

Q2 Where does unit and pricing data come from?
Facility feed, scraping, or manual. The biggest driver of how long this takes.

Q3 Is the existing storage environment the same project?
If so, merge before either goes further.

Q4 Full replacement, or ranked lists alongside WordPress?
We can add ranked pages to the current site as a smaller first step instead of replacing it.
Lower risk, faster to something visible.

Starting today regardless: backups on the current site. That one carries risk every day it waits.


How the engine is put together

For the engineering side of the room — this is the plumbing.ca stack storage.ca would join.

SERVING



BUILDING THE LISTS – runs on a schedule, nobody watching



The important property: pages are built ahead of time and served as files. That's why the engine carries 19 lists without extra cost, and why adding storage is mostly configuration.

Specifics, for the record

plumbing.ca stack

COMPUTE	ECS Fargate ×2, 0.5 vCPU / 1 GB
ENTRY	NLB :443 → :4383
PAGES	S3 + CloudFront, 270 objects
PIPELINE	SQS → 3 Lambdas
DATA	3 DynamoDB tables
SCORED	2,284 businesses

storage.ca stack

COMPUTE	Lightsail small_3_0
REGION	ca-central-1
STACK	WordPress / Apache
SNAPSHOTS	2, both 2024-01-23
DNS_ZONE	orphaned – not authoritative
CDN	outside this account

Read live from AWS account 490318453712 on 23 Sep 2026 – Route 53, ELB, ECS, Lambda, DynamoDB, S3, CloudFront, Lightsail, Cost Explorer. Page counts from storage.ca's own sitemap. A fuller technical walkthrough exists for the engineering review.