

River Refugium Project (RRP)

CERNUNNOS FOUNDATION
BRIGHT MEADOW GROUP
Systems Analysis and Solutions Consulting

Document Title:

RRP6 Economic & Deployment Model

RRP Document Number:

RRP0002.7

Version by date time group:

01 APR 2026 0001

Cernunnos Foundation / Bright Meadow Group

Email: robert@brightmeadowgroup.com

Website: www.cernunnosfoundation.com

****RIVER REFUGIUM PROJECT****

Cernunnos Foundation / Bright Meadow Group

****RRP6 – Economic & Deployment Model****

Document Number: RRP0002.7 | Version: 2.0 | Date: March 2026

Status: Working Draft – Figures Provisional | Reviewed by: [Pending]

| ****Scope of This Document**** RRP6 presents the integrated economic and deployment logic of the River Refugium Project across three scales: pilot, reference node, and watershed arc. It introduces the two deployment models (Model A: integrated on-site HTC/HTL; Model B: distributed biological node feeding a regional hub), the hub-and-spoke cluster architecture, and the modular slot framework for financial modeling under conditions of incomplete crop yield data. *For the detailed revenue layer breakdown and investor-facing financial architecture, see Appendix E – Revenue Stack & Investor Framework.* |
| --- |

****1. The Economic Thesis: Waste Is Feedstock****

The River Refugium Project does not pay to remove pollution. It monetizes it.

Every input that drives environmental damage in a river system – nitrogen, phosphorus, suspended organics, heat, CO₂ – is redirected inside the RRP node and converted into biomass, fuel precursors, carbon-negative materials, and verifiable environmental credits. The pollutant is the feedstock. The cleanup is the production process.

This principle governs every financial decision in the system. It means operating costs are offset by outputs at every stage, and it means the system becomes more economically productive in direct proportion to how polluted the source water is. The most degraded rivers are the best business cases.

This document examines that economic logic at three scales – pilot, node, and watershed arc – and introduces the two deployment models and the cluster architecture that make it financeable in phases.

****2. Three-Tier Economic Scale****

The RRP economic model operates across three distinct scales. Each tier has its own purpose, capital profile, and strategic function. Understanding all three is required to understand any one of them.

| ****TIER 1**** *Biological Pilot* | ****TIER 2**** *100-Acre Reference Node* | ****TIER 3**** *Watershed Deployment Arc* |
| --- | --- | --- |
| 1 full grow unit | 12-14 grow units | Multiple nodes |
| ~7-8 acres | ~100 acres | 2-4 per watershed (typical) |
| 13 greenhouses | Centralized HTC/HTL | Hub-and-spoke clusters |
| 168,000 GPD (est.) | ~2M+ GPD (est.) | Nutrient-load constrained |
| Off-site HTC/HTL | Full revenue stack | Regional industry |
| *Prove the biology. Generate the data that sizes everything downstream.* |
| *Prove the economics. First commercially self-sustaining installation.* | *Scale until the watershed is clean. Market grows as water improves.* |

****Figure 1. ***Three-Tier Economic Scale – Pilot, Reference Node, and Watershed Deployment Arc****

****2.1 Tier 1 – The Biological Pilot (One Full Grow Unit)****

The pilot is not a token demonstration. It is a full-scale biological engine – every functional component of an RRP node operating together at the smallest

economically meaningful unit size. It contains a complete six-tank biofiltration train, one evaporation greenhouse, a 13-house greenhouse complex, a nutrient routing grid, and SCADA/PLC automation.

The pilot is not designed to be independently profitable. It is designed to do one thing above all others:

| **The pilot is the data engine that sizes every reactor cluster that follows.** **No prior dataset exists for an integrated multi-species hydroponic nutrient-processing system at this scale. The pilot generates that dataset. Until it does, all downstream financial models use provisional academic proxies – clearly flagged as such.** |
| --- |

A smaller pilot would generate misleading data. A larger pilot would demand premature capital. One full grow unit is the minimum unit that simulates the full operational complexity of an RRP node, generates crop-specific yield and nutrient-uptake data across real seasonal variation, and produces the biomass flow rates the economic model requires.

The pilot positions the project for grants, credits, and institutional financing – and provides regulators with verifiable water-quality data before any full-scale deployment is proposed.

2.2 Tier 2 – The 100-Acre Reference Node

The 100-acre reference node is the minimum economically self-sustaining installation. It contains 12-14 full grow units feeding a single centralized HTC/HTL conversion block, with full thermal loop integration, continuous biomass throughput, and a complete eight-layer revenue stack.

At this scale, the system crosses from proof-of-function into proof-of-economics. Biomass throughput becomes continuous enough to run HTC/HTL reactors at steady-state. Energy integration loops become efficient. Fixed capital is distributed across enough acreage to reduce per-unit cost. The node operates as a standalone regenerative industrial park.

| **Land Use Factor** | **Per Grow Unit** | **Model A – 100-Acre Node (On-site HTC/HTL)** | **Model B – 100-Acre Node (No On-site HTC/HTL)** |
| --- | --- | --- | --- |
| **Greenhouse footprint (production)** | ~3 acres | ~42 acres (14 units) | ~45-48 acres (15-16 units) |
| **Evap greenhouse + biofiltration + intake** | ~0.6 acres (20% of GH footprint) | ~8-9 acres | ~9-10 acres |
| **Roads + service corridors** | ~0.5 acres | ~7 acres | ~7 acres |
| **Biomass preparation module** | ~0.75 acres | ~10 acres | ~10 acres |
| **HTC/HTL plant + buffer zone** | Off-site (Model B) | ~8 acres (centralized, on-site) | Not applicable – space returned to production |
| **Remaining buffer / expansion** | – | ~23 acres | ~18-20 acres |
| **Total grow units / node footprint** | **~5 acres per unit** | **~14 grow units** | **~100 acres** | **~15-16 grow units** | **~100 acres** |
| **Model B advantage:** **Removing the on-site HTC/HTL plant frees approximately 8 acres, allowing 1-2 additional grow units on the same 100-acre footprint. This increases biomass production capacity and nutrient removal per acre while reducing entry capex. Harvested biomass is transported to a regional Model A hub for processing.** |
| **All figures are design estimates pending pilot calibration. Evaporation greenhouse + biofiltration + intake is modeled at 20% of greenhouse footprint per grow unit. Final land allocation will be refined following pilot site survey and engineering drawings.** |

Table 1. Land allocation per grow unit and by deployment model (Model A: on-site HTC/HTL; Model B: no on-site HTC/HTL). All figures are design estimates pending pilot calibration.

2.3 Tier 3 – The Watershed Deployment Arc

A watershed contains a finite nutrient load. RRP nodes consume that load to produce biomass. Deployment is therefore constrained not by available land, but by available pollution.

This is counterintuitive but important: the system becomes less profitable as the river gets cleaner. That is the correct outcome. It means the financial model is self-limiting in the right direction – expansion stops when the ecological mission is accomplished.

| Node | N Removal (est.) | P Removal (est.) | Biomass Yield | Economics | Decision |
|-------------|-------------------------|-------------------------|----------------------|--|-------------------------------|
| Node 1 | 10-20% | 15-30% | Peak | Strongest returns | BUILD |
| Node 2 | 8-16% | 12-24% | High | Strong returns | BUILD |
| Node 3 | 6-12% | 9-18% | Moderate | Good returns | BUILD |
| Node 4 | 4-8% | 6-12% | Declining | Marginal – incentive dependent | EVALUATE |
| Node 5+ | <4% | <6% | Low | Negative without strong credit framework | MOVE TO NEXT WATERSHED |

Figure 4. Watershed Saturation Curve – Nutrient removal, biomass yield, and economic returns across successive nodes in a single watershed. All removal percentages are estimates pending pilot data.*

For a typical 100,000-acre mixed agricultural watershed, most basins support 2-4 economically viable nodes. High-nutrient-density corridors such as the Lower Mississippi can support more. The practical approach is to build until the marginal return on the next node no longer justifies the capital, then move to the next watershed.

This is not failure. It is the definition of success: the river is clean enough that further nodes are unnecessary.

**3. Two Deployment Models

The RRP can be built in two configurations depending on site scale, capital availability, and regional infrastructure context. These are not competing designs – they are complementary roles in a regional cluster architecture.

| | MODEL A – Integrated Node | MODEL B – Distributed Biological Node |
|--------------------------|---|--|
| HTC/HTL Plant | On-site, centralized | Off-site regional hub |
| Entry Capex | Higher – full reactor required at build | Lower – reactor investment deferred or shared |
| Energy Loop | Full closed loop: waste heat → greenhouses | No on-site heat recovery; transport adds cost |
| Biomass Handling | Slurry piped directly to reactor | Baled / slurried, transported to hub |
| Best For | 100-acre reference nodes, peak nutrient sites | Pilots, constrained sites, hub satellite nodes |
| Maintenance Risk | Single-site downtime affects full system | Hub fed by multiple satellites – resilient |
| Cluster Role | The hub in a hub-and-spoke cluster | The spoke – feeds a shared regional hub |
| Financial Profile | Higher capex, higher integrated returns | Lower entry cost, revenue share with hub |
| Pilot Suitability | No – too capital-heavy for proof stage | Yes – correct scale for data collection |

Figure 2. Model A vs. Model B – Deployment configuration comparison across key engineering and financial dimensions.*

3.1 Model A – Integrated Node (On-Site HTC/HTL)

Model A is the full closed-loop configuration. The HTC/HTL thermochemical plant sits on-site alongside the greenhouse complex. Biomass moves from harvest to reactor by pipe or conveyor with no transport cost. Waste heat from the reactor feeds back into the greenhouses, reducing energy costs. The aqueous phase from thermochemical processing recirculates into the algae production loop.

Model A is the target configuration for 100-acre reference nodes at high-nutrient sites where biomass throughput is sufficient to run the reactor at steady-state. It is capital-intensive at entry but delivers the highest integrated return profile.

Model A is also the hub in a hub-and-spoke cluster. It anchors a regional group of Model B satellite nodes, receiving their biomass output and processing it through its reactor. The hub's economic viability depends on the aggregate biomass supply from the cluster – which is why cluster sizing is the critical design variable.

3.2 Model B – Distributed Biological Node (Off-Site HTC/HTL)

Model B deploys the full biological engine – biofiltration, evaporation greenhouse, greenhouse complex, routing grid – without an on-site thermochemical plant. Harvested biomass is baled, slurried, or pelletized and transported to a regional Model A hub for processing.

Model B has a significantly lower entry capital requirement. It is the correct configuration for the pilot, for constrained or smaller sites, and for the satellite nodes in a hub-and-spoke cluster. It can begin generating nutrient removal data and biomass revenue before a regional hub exists – with biomass initially sold to third-party processors while the hub is under development.

Model B is not a lesser version of Model A. It is the deployment instrument that allows the network to grow ahead of reactor infrastructure, and the configuration that makes phased regional buildout financially feasible.

4. The Hub-and-Spoke Cluster Architecture

The fundamental operational challenge of any industrial reactor is feedstock continuity. HTC/HTL equipment is capital-intensive, has defined maintenance cycles, and runs most efficiently at steady-state. A single biological node cannot guarantee that continuity – crop rotation gaps, seasonal variation, disease quarantine, and storm events all create supply interruptions.

The hub-and-spoke architecture solves this. One Model A node anchors a regional cluster as the processing hub. Between five and ten Model B biological satellite nodes supply it with biomass from different sites, different crop profiles, and different seasonal curves. When one satellite has a gap, the others are producing. The hub reactor stays fed. Maintenance windows can be scheduled against the aggregate supply curve rather than a single site's output.

```
| **REGIONAL HTC/HTL HUB (Model A Node)** *Reactor operates at steady-state fed  
by aggregate cluster biomass* |  
| --- |  
| ↑ Biomass delivery | ↑ Biomass delivery | ↑ Biomass delivery |  
| **Satellite Node 1** ~7-8 acres Full grow unit Model B | **Satellite Node 2**  
~7-8 acres Full grow unit Model B | **Satellite Node 3** ~7-8 acres Full grow  
unit Model B |  
| *+ 2 to 7 additional satellite nodes (cluster size determined by pilot data)*  
|  
| **CLUSTER LOGIC** Reactor minimum feed requirement ÷ conservative biomass  
yield per node = minimum cluster size *Actual cluster size determined by pilot  
data. Projected range: 1 hub + 5-10 satellites for typical nutrient-dense
```

watershed.* |

Figure 3. Hub-and-Spoke Cluster Architecture – One Model A integrated hub node anchored by 5–10 Model B biological satellite nodes. Cluster size determined by pilot data.*

4.1 The Cluster Sizing Problem – And Why the Pilot Solves It

The correct cluster size – how many satellite nodes are needed to keep the hub reactor running at minimum efficient throughput – depends on one number: how much biomass a single grow unit produces per unit time, consistently, across seasons.

That number does not exist yet. No prior system has run a multi-species hydroponic nutrient-processing network at this scale under real-world conditions. Academic yield data for individual species exists, but the integrated system performance – accounting for nutrient variability, seasonal rotation, disease events, and maintenance cycles – has never been measured.

This is not a gap in the design. It is the founding justification for the pilot.

The pilot measures exactly this. Over its operational life it generates the biomass yield curves, the seasonal variation data, the maintenance cycle impacts, and the crop diversity resilience data that the cluster sizing model requires. Until those numbers exist, all cluster size estimates are modeled from academic proxies. The range of 5–10 satellites per hub reflects that uncertainty honestly – it will narrow as pilot data accumulates.

4.2 Why This Is Financially Honest and Fundable

The temptation in early-stage infrastructure proposals is to project with false precision. To pick a crop yield number, multiply it across acreage, and present a clean pro forma that implies certainty that does not exist.

This proposal takes the opposite position. The financial model is explicitly modular – built on slot assumptions rather than species assumptions, with provisional figures clearly flagged, and with the pilot framed as the instrument that replaces those figures with real data.

This is more fundable, not less. Sophisticated capital partners and agency program officers have seen enough inflated projections to distrust them. A model that says 'here is what we know, here is what we are measuring, and here is how the numbers update when we have data' is a model they can engage with.

5. The Modular Slot Framework

Because actual per-species yield data for this integrated system does not yet exist, the financial model is built on slot assumptions rather than crop assumptions. Each greenhouse in the RRP complex is treated as a standardized production slot with defined characteristics – turnover cycle, nutrient uptake profile, biomass output range, and HTC/HTL suitability – regardless of which specific species fills it.

When pilot data becomes available, it updates the slot parameters. The financial model recalculates automatically. The structure is consistent across all projections; only the input numbers change as evidence accumulates.

| Slot Type | Growth Profile | Cycle Time | Nutrient Uptake | HTC Suitability | Data Source |
|------------------|-----------------------|-------------------|------------------------|------------------------|--------------------|
|------------------|-----------------------|-------------------|------------------------|------------------------|--------------------|

| | | | | | |
|-----|-----|-----|-----|-----|-----|
| --- | --- | --- | --- | --- | --- |
|-----|-----|-----|-----|-----|-----|

| | | | | | |
|--------------------|-------------------------------|---------------|------------------------|-----------|--|
| Slot Type A | High-turnover aquatic / algae | Days to weeks | Very high N + P uptake | Excellent | Chlorella, Spirulina literature (academic) |
| PROVISIONAL | | | | | |

| | | | | | |
|--------------------|------------------------|------------|--------------------|--|--|
| Slot Type B | Fast-fiber terrestrial | 6–12 weeks | High N, moderate P | | |
|--------------------|------------------------|------------|--------------------|--|--|

Good | Hemp, miscanthus academic yield data **▮ PROVISIONAL** |
| **Slot Type C** | Slow-fiber / structural | 12-24 weeks | Moderate N + P |
Good – high lignin | Willow, bamboo, jute literature **▮ PROVISIONAL** |
| **Slot Type D** | Textile fiber | 12-20 weeks | Moderate N, low P | Moderate |
Cotton, flax academic data **▮ PROVISIONAL** |
| **Slot Type E** | Aquatic / wetland macrophyte | 4-8 weeks | High P, moderate
N | Good | Cattail, duckweed academic data **▮ PROVISIONAL** |
| **▮ PROVISIONAL: ***All yield and uptake figures in this table are derived
from published academic literature for individual species under controlled
conditions. They are placeholders only. Actual per-slot performance data will be
generated by the RRP pilot and will replace these proxies. See RRP8 for the
pilot data collection protocol.* |

Figure 5. *Modular Slot Framework – Greenhouse slot types with provisional
academic yield proxies. All figures flagged PROVISIONAL pending pilot data
collection per RRP8.*

5.1 How the Framework Is Used

Each of the 13 greenhouses in a grow unit is assigned a slot type based on its
physical configuration, nutrient routing position, and current crop assignment.
The financial model then calculates:

- Biomass output per slot per grow cycle (provisional range from literature)
- Nutrient uptake per slot per cycle (provisional – determines credit
eligibility)
- HTC/HTL feedstock contribution per slot (provisional – determines reactor load
planning)
- Annual aggregate output per 13-slot grow unit (sum of slot outputs across the
year)
- Cluster-level aggregate output (sum across all satellite nodes in the hub
cluster)

This framework allows consistent financial modeling across all three tiers and
both deployment models without requiring crop-specific certainty that does not
yet exist.

5.2 The Provisional Data Protocol

All provisional figures in this framework carry a flag in tables and
projections. They are sourced from peer-reviewed academic literature for
individual species under controlled conditions – the closest available proxy for
RRP slot performance.

The RRP8 Verification & Monitoring document defines the pilot data collection
protocol that will replace these proxies with measured values. As pilot data
accumulates, updated figures will be published in the open-access RRP data
repository and incorporated into revised versions of this document.

The financial model is designed to be updated. The structure does not change.
The numbers do.

6. Strategic Summary

The RRP economic model is governed by three structural truths:

| **1** | The pilot is not a demonstration – it is the data collection
instrument that sizes every reactor cluster that follows. Its value is not in
what it produces but in what it measures. |
| --- | --- |

| ****2**** | The 100-acre reference node is the minimum economic unit. Below this scale, thermochemical conversion is not commercially viable. Above it, returns improve as fixed costs spread across more acreage. |
| ****3**** | The hub-and-spoke cluster is the operating model for regional scale. One Model A hub anchored by 5-10 Model B satellites is the expected norm. Cluster size is determined by pilot data, not desktop projection. |

The financial model is intentionally modular. Slot assumptions replace species assumptions. Provisional figures are flagged. Real data replaces them as the pilot generates it. This is not a limitation of the model – it is its most defensible feature.

For the detailed revenue layer breakdown, capital stack analysis, and investor-facing financial architecture, see Appendix E – Revenue Stack & Investor Framework.

****References****

1. Bright Meadow Group, River Refugium Project Technical Whitepaper. CF/BMG Internal, 2025.
2. RRP Biological Unit Design Notes, CF/BMG Systems Draft, 2025.
3. Biomass Yield & Routing Grid Validation Plan, RRP Pilot Site Framework, 2025.
4. Grow Unit Footprint and Land Allocation Model, RRP Site Planning Draft, 2025.
5. Watershed Nutrient Removal Modeling, RRP Technical Estimates §12.1-12.4, 2025.
6. USGS/USDA Nutrient Export Data for Mixed Agricultural Basins. National print editions.
- *7. Academic yield proxies: Chlorella/Spirulina – Mata et al. (2010); Hemp – van der Werf (2004); Miscanthus – Heaton et al. (2008); Willow – Haughton et al. (2009); Duckweed – Körner et al. (2003). [Note: These are placeholders; full citations to be standardized in final version.]*