

RIVER REFUGIUM PROJECT

Watershed-Scale Nutrient Recovery Infrastructure

APPENDIX A

Site Selection Matrix

Ten weighted criteria for evaluating candidate node locations

Document No. RRP0003-apA | Version 3.0 | August 2026

Companion to RRP0003 — Consolidated Technical Whitepaper, Version 3.0

Supersedes RRP0002-AppA, Version 2.0 (01 April 2026)

CERNUNNOS FOUNDATION

BRIGHT MEADOW GROUP

robert@brightmeadowgroup.com | www.cernunnosfoundation.com/rrp

Open Access | No Paywall | Reproduction and adaptation permitted

ABSTRACT

This appendix defines the standardized scoring tool used to evaluate candidate locations for River Refugium Project pilots, 100-acre nodes, regional clusters, and long-term watershed deployment. Each site is scored across ten weighted criteria on a 0–5 scale, producing a composite out of 150.

The matrix weights financial and engineering viability. Demonstrated environmental need is not itself a scoring criterion, because a location that cannot support a solvent facility does not become able to by deserving one. Need determines whether the framework should exist; the matrix determines where it can first be built successfully.

The recommended minimum threshold for first-wave nodes is 110 of 150. A worked example using the Helena-West Helena / Lula corridor, the top-ranked site in Appendix C, is included as a demonstration of the arithmetic.

Version 3.0 note: the scope of this matrix is stated explicitly in Section 3 and is narrower than Version 2.0 implied. It is calibrated for main nodes and initial test sites, not as a general map of where the framework has application.

1. Purpose and Structure

Site selection for a facility of this type fails in one of two directions. Chosen on environmental need alone, the first node is built where the water is worst and the economics are hardest, and it fails financially before it produces a usable dataset. Chosen on economics alone, it is built where nothing much needed fixing and it proves nothing.

The matrix is constructed to hold both. Nutrient load carries the heaviest weight because it is simultaneously the environmental case and the feedstock supply. The remaining criteria establish whether a facility placed against that load can actually be built, staffed, permitted, and operated.

Each criterion is scored 0–5 and multiplied by its weight. Weights total 25, so the maximum composite is 150.

2. Scoring Criteria and Weights

#	Criterion	What Is Being Assessed	Weight	Helena Score
1	Nutrient Load & Water Quality	Basin nutrient density, turbidity, dissolved and suspended organics	×4	5/5 = 20
2	Land Cost & Availability	Acreage cost, zoning, floodplain constraint, assembly difficulty	×3	5/5 = 15
3	Climate Compatibility	Growing season length, freeze risk, solar availability for evaporation stage	×3	5/5 = 15
4	Hydrological Stability	Flow reliability, drought exposure, surge and flood-stage risk	×3	5/5 = 15

#	Criterion	What Is Being Assessed	Weight	Helena Score
5	Labor Availability	Skilled and semi-skilled workforce access; training infrastructure	×2	4/5 = 8
6	Regulatory Environment	Permitting difficulty, agency posture, compliance friction	×2	4/5 = 8
7	Logistics & Access	Road and rail, utilities, distance to buyers and to hub reactor	×2	4/5 = 8
8	Local Investment Alignment	County, state, and tribal willingness to participate in capital	×2	5/5 = 10
9	Community Compatibility	Local support, siting history, risk of organized opposition	×1	4/5 = 4
10	Strategic Expansion Value	Influence on regional scaling and on subsequent site viability	×3	5/5 = 15

Helena-West Helena / Lula composite: 118 of 150. This exceeds the 110 first-wave threshold and is ranked flagship deployment site in Appendix C.

3. Scope of the Matrix

The matrix is calibrated for main nodes and initial test sites — locations where nutrient load, hydraulic access, land availability, and logistics are strong enough to support a full facility and, more importantly, to generate a clean dataset quickly.

It is a filter for where to build first. It is not a map of where the framework has application, and a low composite score is not a finding that the framework cannot serve a location. Many sites that score poorly here — thinner nutrient loads, single dominant contaminants, small municipal or industrial discharge problems — are legitimate targets for configurations that cannot yet be specified, because the data required to size them does not exist until the first nodes run.

Why need is not weighted

A weighted need criterion would be double-counting. Nutrient load already carries the heaviest weight in the matrix, and nutrient load is the measurable form of need at watershed scale.

Adding a separate need score on top of it would allow a site to be selected on sympathy after the physical case had already been counted, which is the failure mode that produces demonstration projects that close when the grant ends.

4. Scoring Discipline

- Score against evidence, not intention. A county that says it will help scores on what it has done for comparable facilities, not on what it says at a meeting.
- Score criterion 4 against the flood record, not the average. A site with excellent mean flow and a history of stage events that would overtop a forebay is a low score, not a moderate one.

- Score criterion 9 honestly and low when uncertain. A community with a history of being sited upon is a real constraint on a facility of this footprint, and treating it as a formality is how projects lose years late rather than weeks early.
- Do not adjust a score after a composite is calculated. If the composite is disappointing, that is information about the site.

Scores represent relative deployment readiness. They are not performance guarantees, and site-specific survey is required before any capital commitment.

5. Open Question for Version 3.0

The matrix does not currently account for which production configurations a site's water chemistry actually permits. A site with strong hydraulic and logistical scores whose chemistry supports only the sequestration configuration has a materially different revenue profile than its composite suggests, because sequestration output is the least certain revenue class in the model.

Two candidate remedies are under consideration: an eleventh criterion scoring configuration breadth, or a separate overlay applied after scoring that flags composite scores whose revenue basis is narrow. The second is preferred at present because it does not disturb the existing weights or invalidate scores already published.

■

Notes

Citations follow Chicago Notes-Bibliography style. Internal Cernunnos Foundation and Bright Meadow Group documents are cited by title and year. Figures marked ■ are provisional academic proxies pending replacement by measured pilot values.

1. Bright Meadow Group, "RRP Site Selection Methodology." CF/BMG Internal, 2025.
2. Bright Meadow Group, "RRP0003-apC — Initial Five Recommended Sites," Version 3.0. CF/BMG, August 2026.
3. Bright Meadow Group, "RRP0003 — Consolidated Technical Whitepaper," Version 3.0, §14.2. CF/BMG, August 2026.
4. USGS, SPARROW Model — Nutrient Load Estimates for Major U.S. Basins. National Water-Quality Assessment Program.