VENDOR EVALUATION MEMO · LUMEN HEALTH

AI Vendor Evaluation Report - ridgeline_stroke_lvo

Generated: 2026-04-16T14:32:00Z **Vendors evaluated:** neuralstroke, omnirad
Rubrics applied: governance, safety, implementation, auditability

1. Executive Summary

Recommendation: phased selection of NeuralStroke, contingent on three deliverables within 90 days. NeuralStroke has the stronger clinical evidence base for stroke LVO detection, more mature post-market surveillance, and a better-characterized rural transfer workflow than OmniRad. OmniRad offers consolidation economics and a broader module footprint, but its stroke-specific evidence and community-hospital operational track record are thinner and cannot be reconciled within the decision timeline. The committee should not sign the unconditional contract: two material gaps remain — disaggregated demographic performance for Native American and rural Medicare subgroups, and a tested multi-PACS rollback at the Fuji Synapse community sites. One disqualifying finding stands against OmniRad on auditability (`logging_and_event_capture`, score 2) driven by a default retention period under the Ridgeline clinical record policy. The biggest risk is community-site alert fatigue; Ridgeline's Aidoc experience is directly relevant and should shape the go-live gates.

2. Case Context

Ridgeline Health is a 9-hospital system anchored by a Level I / comprehensive stroke center in Denver, with three primary stroke centers and two rural hospitals more than 90 minutes from thrombectomy capability. The stroke medical director has proposed AI-assisted LVO detection to compress door-to-needle and door-to-groin-puncture intervals across the system. Epic is deployed system-wide; imaging runs on Sectra at the flagship and Fuji Synapse at community sites. The committee's standing priorities — evidence grounded in the deployment population, auditability sufficient for regulatory defense, and a credible implementation path at non-flagship sites — are the frame for this review.

3. Scorecards

3.1. Governance

CRITERION	NEURALSTROKE	OMNIRAD	NOTES
`regulatory_status`	4	3	NS has PCCP submitted; OR narrower indication for community sites.
`contractual_terms`	3	3	Both: uptime SLA + audit rights; neither contracts to clinical performance floor.
`vendor_stability`	3	4	OR has broader customer base and a documented continuity plan.
`policy_alignment`	4	3	NS publishes HTI-1-conformant model card; OR has it in preparation.
`liability_allocation`	3	3	Both cap at 2x annual fees; neither includes bias indemnity by default.
`privacy_and_data_use`	4	3	NS requires explicit opt-in for secondary use; OR defaults to opt-out.
Weighted score	**3.57**	**3.20**	
Qualified	yes	yes	

Top NeuralStroke finding: `regulatory_status` (4) — PCCP submitted, cleared indications match flagship and community deployment; rural transfer language is implied but not explicit. Missing: the clearance letter and PCCP approval status. Top OmniRad finding: `vendor_stability` (4) — diversified customer base and tested continuity procedures outweigh the narrower stroke-specific evidence for this rubric.

3.2. Safety

CRITERION	NEURALSTROKE	OMNIRAD	NOTES
`clinical_evidence_quality`	4	3	NS has three multi-center studies; OR has one prospective plus two retrospective.
`bias_and_equity_testing`	3	2	Neither has Native American subgroup data; OR will not commit to deliver.
`failure_mode_characterization`	4	3	NS publishes a quantitative failure mode catalog; OR offers narrative only.
`post_market_surveillance`	4	3	NS provides site-level monitoring data; OR aggregate only.
`alert_and_output_burden`	3	2	OR override rates in reference-check trend high; concerning given Aidoc history.
Weighted score	**3.58**	**2.62**	
Qualified	yes	yes	NS threshold met; OR just clears `bias_and_equity_testing` floor at 2.

Top NeuralStroke finding: `clinical_evidence_quality` (4) — multi-center prospective evidence with demographic diversity approaching Ridgeline's mix, though Wind River Native American representation remains a gap (missing_information). Top OmniRad finding: `alert_and_output_burden` (2) — override rate data from comparable community deployments is directionally worse than NeuralStroke's and directly relevant given Ridgeline's Aidoc ICH experience.

3.3. Implementation

CRITERION	NEURALSTROKE	OMNIRAD	NOTES
`ehr_integration_architecture`	4	4	Both have production multi-PACS Epic deployments.
`workflow_fit`	4	3	NS has published operational results at comparable multi-tier systems.
`training_and_change_management`	3	4	OR provides an embedded clinical transformation specialist.
`it_burden_and_infra_fit`	4	3	OR's rural bandwidth requirements are not validated at Ridgeline's sites.
`migration_and_rollback`	3	3	Neither supports tested per-site parallel-run with Aidoc.
`tco_realism`	3	4	OR's TCO is more interrogable; NS omits incumbent decommissioning.
Weighted score	**3.55**	**3.48**	
Qualified	yes	yes	

Top NeuralStroke finding: `workflow_fit` (4) — documented rural transfer patterns and ED overread configurations map to Ridgeline's community sites. Top OmniRad finding: `training_and_change_management` (4) — embedded transformation specialist is a real implementation advantage not matched by NeuralStroke.

3.4. Auditability

CRITERION	NEURALSTROKE	OMNIRAD	NOTES
`model_transparency`	4	3	NS model card is clinical-grade; OR is data-science-grade.
`logging_and_event_capture`	4	2	**OR default retention 18 months; Ridgeline policy requires 7 years.**
`monitoring_and_drift_detection`	3	3	Neither offers contractual notification-latency bound.
`incident_response_readiness`	3	2	OR has not supported a customer MDR filing; NS has one on record.
`data_lineage_and_retention`	4	3	NS provides end-to-end lineage with model-version tie-in.
Weighted score	**3.64**	**2.58**	
Qualified	yes	**no**	OR disqualifying: `logging_and_event_capture`

Top NeuralStroke finding: `logging_and_event_capture` (4) — customer-accessible logs with user, version, and clinician action captured; retention configurable to 7+ years. Top OmniRad finding: `logging_and_event_capture` (2, disqualifying) — default retention is below Ridgeline's clinical record policy and the vendor will not commit contractually to a configurable 7-year retention without a custom engagement.

4. Synthesis

Averaged weighted scores favor NeuralStroke on every rubric, but the decision is not an average. OmniRad's disqualification on auditability is the dominant signal: an AI tool that cannot be reconstructed during an adverse event review creates unbounded liability for the operator. Committee priorities — rural transfer performance, alert fatigue avoidance given Aidoc history, and regulatory defensibility — all line up with NeuralStroke's profile. The strongest counter-argument is consolidation: OmniRad brings a module roadmap that NeuralStroke cannot match, and the committee is giving up future vendor leverage by choosing a single-use platform. That tradeoff is real but not dispositive in Q2 2026.

5. Recommendation

Phased selection of NeuralStroke (vendor_a), conditional on the deliverables in Section 6. NeuralStroke's evidence base, post-market surveillance maturity, and auditability posture meet the committee's bar; OmniRad's auditability gap is disqualifying as presented. A phased deployment — flagship plus two primary stroke centers at 6 months, full community-site expansion after defined performance gates — manages the community-site and rural subgroup risks that neither vendor has fully answered.

6. Conditions

1. NeuralStroke delivers disaggregated demographic performance for Native American and rural Medicare subgroups within 90 days of signing; if not delivered, Ridgeline exercises termination without penalty.
2. NeuralStroke provides the approved FDA PCCP letter and a written mapping of algorithm update types to customer notification commitments within 60 days of signing.
3. NeuralStroke agrees to a contractual notification-latency bound for drift detection (7 days or less) and to provide site-level monitoring data access for Ridgeline's compliance team.
4. Community-site go-live is gated on Aidoc PE module performance review and stroke coordinator workflow validation at Wyoming and Kansas hospitals.
5. A contractual off-ramp permits Ridgeline to revert to an Aidoc-only configuration at any community site during a 12-month stabilization window.

7. Tradeoffs Accepted

- Single-use stroke platform rather than a consolidated multi-module radiology AI footprint; future negotiating leverage with a broad-platform vendor is reduced.
- Smaller vendor customer base than OmniRad; elevated acquihire and continuity risk managed via source code escrow.
- Loss of OmniRad's embedded clinical transformation specialist; Ridgeline must fund in-house change management for community sites.

8. Risks and Mitigations

RISK	LIKELIHOOD	IMPACT	MITIGATION
Model underperformance in Wind River Native American subgroup	moderate	high	Contractual subgroup monitoring and termination trigger.
Alert fatigue at community sites mirroring Aidoc ICH experience	moderate	high	Alert-threshold configuration by role and time-of-day; go-live gate on override-rate thresholds.
NeuralStroke customer concentration / acquihire risk	moderate	high	Source code escrow with technical trustee; assignment protections.
Fuji Synapse PACS integration scope underestimated	moderate	moderate	Vendor references at multi-PACS systems verified before signing; per-site rollback documented.
PCCP scope uncertainty — future algorithm updates may require new 510(k)	low	moderate	Change-notification commitments in contract; customer-side validation gate for major updates.
Vendor will not contract to bias indemnity	moderate	moderate	Separate cyber + product liability coverage with customer named; legal review of exclusions.
Aidoc renewal posture shifts once Ridgeline commits to NeuralStroke	low	moderate	Decouple contracts; maintain Aidoc PE module through stabilization window.

9. Open Questions

- **Governance:** FDA 510(k) clearance letter; PCCP approval letter; subprocessor list at signing; cyber and product liability policy limits.
- **Safety:** Native American and rural Medicare subgroup performance; community-hospital override-rate benchmarks; non-fellowship reader overread performance data.
- **Implementation:** Fuji Synapse reference deployments; rural bandwidth validation at comparable sites; per-site FTE and calendar commitments.

- **Auditability:** Contractual notification latency for drift; MDR filing support letters; end-to-end lineage demonstration at a reference site.

10. Appendix - Methodology Notes

Rubric versions used: governance v0.1, safety v0.1, implementation v0.1, auditability v0.1. Weighted scoring uses the published per-criterion weights; no per-case weight overrides were applied. Qualification follows the `minimum_threshold` logic defined in each rubric — any finding below its threshold triggers a disqualifying classification regardless of weighted score. Evidence cutoff is the case materials as of 2026-04-16.

Published as part of Lumen’s open vendor-evaluation framework.
Web: lumenhealth.ai/vendor-evaluation · Contact: hello@lumenhealth.ai