

7 September 2026

Reducing Emissions from Deforestation and Forest Degradation (REDD) FRC Feedback

-Responses to Stakeholder Questions and Comments-

The following document presents Sylvera's responses to questions and feedback submitted during the Framework Review Committee (FRC) for comments period on the Sylvera REDD Framework Version 4.0.

CARBON ACCOUNTING

1. Carbon Accounting and Equitable Earth (EQE)

ARCA is currently developing projects exclusively under the Equitable Earth standard, meaning all carbon accounting calculations are managed directly by the standard (EQE). Given this structure, how will Sylvera's evaluation and reporting process approach the optimization of calculations and the mitigation of any methodological limitations identified? Furthermore, are there mechanisms to prevent adverse scoring impacts (penalties) that might arise from perceived vulnerabilities in the EQE methodology when the Carbon Accounting Assessment?

Response:

Sylvera understands that some calculations in the M002 Methodology for Terrestrial Forest Conservation are performed directly by the Equitable Earth (EE) standard rather than by the project developer. This was considered as we built our framework. While Sylvera assesses projects in a methodology-agnostic way, we stay aware of what a project can and cannot choose through rating commentary.

We previously assessed this methodology in a methodology profile and found its approach to be of high integrity; this means that the standard-calculated sections should score well. However, as always, each project is assessed on its implementation. Therefore, we cannot ignore certain methodological choices, so we have not developed a mitigation for



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

methodological limitations. Below, we break down how we assess sections decided by the Equitable Earth standard versus those decided by the developer, and we hope to demonstrate the framework we use to address the M002 methodology.

The standard determines the baseline. This is the largest standard determination, and Sylvera cannot ignore this section; it must evaluate whether it's appropriate, as this is core to assessing a REDD project's integrity. We run a quantitative baseline evaluation using our dynamic baselining approach, which observes the counterfactual rather than projecting it. This bypasses the deforestation risk map and enables more detailed analysis through comparison with ex-post emissions. We also qualitatively assess the deforestation maps used, tied to the choice of map and modeling approach. In tandem with this, we analyze the baseline scenario in the additionality section, as defined by the project proponent.

Equitable Earth calculates a project's carbon stock. We generate a range of independent values and compare them to the project's figures. This ensures projects aren't penalized harshly against a Sylvera-calculated figure. Equitable Earth also measures project emissions. Per the methodology documentation, this relies on remote sensing, an approach that scores highly in our framework; this would not limit the rating. Lastly, uncertainty is also calculated per the methodology. Equitable Earth uses a Monte Carlo simulation to propagate uncertainty scores, which is considered best practice in the framework.

Baselining

2. The counterfactual is never validated against the outcome time series. Covariate balance is demonstrated; predictive performance is not. This is the framework's central gap.

Model validation against the outcome time series.

Comment: The framework demonstrates that project and control pixels are similar on observed covariates at a point in time. It nowhere demonstrates that the matched set reproduces the project's deforestation trajectory and trends. Balance on observables does not establish predictive validity: a well-balanced matched set can still track the



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

outcome poorly, and a model that fits a historical window through overfitting or data leakage can fail entirely in the monitoring period. Slide 27 lists model/placebo error as “exploring,” so we take it this is recognised – but the unvalidated counterfactual is already generating the plausibility range used to score projects on slide 29. In our experience this is not a marginal term. Model error characterised through pseudo-projects is the single largest component of uncertainty in an AUDD dynamic baseline, and the one that cannot be inferred from balance diagnostics.

Ask:

- *Before the plausibility range is used for scoring, report: Pre-intervention error as a percentage of mean annual project deforestation, across the full pre-intervention window;*
- *Error for a held-out validation window not seen by the matching covariates, reported separately – the more diagnostic of the two figures;*
- *Forward validity demonstrated over at least three consecutive years after control selection;*
- *Out-of-sample error quantified on spatially independent pseudo-projects.*

Sylvera Response:

We match at project start year, considering only pixels that have been continuously forested over the ten years prior, so the pre-intervention deforestation trajectory is zero by design. We recognise that the lack of model error quantification via placebo testing is a gap in the current approach; we plan to continue testing the model, including assessing and propagating model error using pseudo-projects, and this will be incorporated into a future version.

3. SMD is used as the primary confidence signal, and it is not an uncertainty metric. Pathway assignment should turn on whether the model reproduces observed outcomes, not on covariate similarity.

SMD as the confidence gate

Comment: Standardised mean difference is a balance diagnostic – a first-moment statistic describing similarity of inputs. It carries no information about error in the



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

estimated outcome. Using it as the primary signal for pathway assignment substitutes a statement about how alike the pixels look for a statement about how wrong the counterfactual might be. Two issues compound this: $SMD \leq 0.4$ is lenient against conventional practice, where 0.1/0.2 is treated as good balance; and SMD alone cannot detect distributional mismatch, since two distributions can share a mean exactly and differ substantially in shape.

Ask:

- *Base pathway assignment on demonstrated predictive performance – validation-period RMSE and pseudo-project out-of-sample error – rather than on covariate balance.*
- *Retain balance as a necessary but not sufficient precondition, tightened to $SMD \leq 0.2$ and supplemented with a distribution-agnostic overlap metric (e.g. PDF overlap ≥ 0.7).*
- *Add population density to the covariate set; it is a primary proximate driver of smallholder clearing and is currently absent.*

Sylvera Response:

Pathway assignment currently still runs on balance rather than predictive performance, since we don't yet have the pseudo-project error metrics to base it on; we agree that's the right long-term target, and once the pseudo-project work lands (see previous point) we intend to shift pathway assignment onto validation-period error rather than balance alone.

In the meantime, we've tightened the balance threshold to $SMD \leq 0.2$ in the current release, in line with your suggestion. We also agree mean-based balance alone can miss distributional mismatch, and in the next release we'll move beyond SMD to adopt further distribution-aware metrics.

On covariates: we currently use proximity to settlements, roads, and deforestation in the five years before the project start date. We agree that population density can be a meaningful driver in some contexts, we'll explore adding it to the covariate set in the next release.



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

4. The plausibility range is under-propagated, and it is the scoring gate. How a project's baseline is compared to it is not stated, so it is unclear whether a developer who propagates uncertainty fully is rewarded for it, ignored, or penalised.

Uncertainty propagation, and the consequence of an over-narrow band

Comment: Slide 27 states that spatial autocorrelation is not modelled and that the band is "likely optimistic." Slide 16 shows it is handled for the Biomass Atlas, via semivariogram analysis, expressly to avoid understating uncertainty in spatially correlated forest environments. The same problem is therefore corrected in one half of the framework and left uncorrected in the half that produces the scoring gate – where matched donors are if anything more spatially clustered, since k-NN matching actively selects neighbouring pixels with similar covariates. The framework then does not state how a project's baseline is compared to the range. Slide 29 shows the project baseline as a single line, which suggests the project's own uncertainty may not enter the comparison at all. Three readings are possible, with very different consequences for a developer who propagates uncertainty fully:

- *Interval against interval. A project reporting a fuller, wider interval overlaps more easily and passes; a project reporting a bare point estimate is judged on that point alone. Rigour is rewarded.*
- *Project central estimate against the band. The developer's uncertainty work is invisible to the score. Rigour earns nothing, and an under-propagated band produces divergence flags that reflect error in Sylvera's estimate rather than in the project's.*
- *Credited volume against the band. Where a developer credits conservatively at the lower bound of their interval rather than the central estimate, the more conservatively they credit, the further below the range they sit. If divergence below the range is scored like divergence above it, conservative crediting is penalised directly.*

We do not believe any of these outcomes is intended, and the first is clearly the right one. But as drafted the framework does not tell a developer which applies – and under the second and third, the developers doing the most rigorous uncertainty work are the ones who gain least from it or are actively disadvantaged.

Ask:



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

- *Extend the semivariogram treatment of spatial autocorrelation to the dynamic baseline.*
- *Include model error, characterised via pseudo-projects, in the uncertainty stack.*
- *Combine model, activity-data and carbon-stock uncertainty into a single Monte Carlo interval on the avoided-emissions estimate.*
- *State how a project baseline is compared to the range: whether the project's own uncertainty interval is used or only its central estimate, and whether the figure compared is the modelled baseline or the credited volume after conservative deductions.*
- *Confirm that a project reporting fuller uncertainty propagation than Sylvera's own is not disadvantaged by doing so.*

Sylvera Response:

We'll introduce spatial autocorrelation into the dynamic baseline in the next release, alongside model error from pseudo-projects. On canopy height specifically, our current uncertainty propagation assumes a fixed match through the annual loss product, which is a simplification we're aware of and plan to revisit. On how a project's baseline is compared to the range: the comparison itself is Sylvera's plausibility range against the project's own baseline, rather than interval against interval. Specifically, we compare against whichever figure the project actually credits off (typically the central estimate) but where a developer credits more conservatively off a lower or upper bound of their own uncertainty range, that bound is what we factor into the comparison instead. A developer's own uncertainty and propagation isn't usually factored into that comparison score directly but it is looked at separately as part of the assessment, and rewarded where appropriate propagation, methods and transparency are applied

5. Pathway logic (A / B / C)

Comment: We support the underlying design – quantitative assessment where it is



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

reliable, a qualitative route where it is not, with variable weighting between them. Pathway assignment determines whether a project's baseline is judged primarily on Sylvera's quantitative comparison or primarily on qualitative review. It is therefore a consequential decision, and a developer should be able to apply the rule to their own project and arrive at the same answer Sylvera does. As specified, they cannot:

- A and B state identical dynamic baseline criteria ($SMD \leq 0.4$, matched $\geq 75\%$), so the split rests entirely on an undefined distinction between "high MSL coverage" and "coverage."
- The boundaries overlap. A project at exactly SMD 0.4, or exactly 75% matched, satisfies both the A/B and C criteria.
- Projects can fall through all three. No AND/OR is stated. A project at SMD 0.3 with 60% match coverage satisfies neither set – a common combination, since balance and coverage are traded against each other by the matching algorithm itself.
- The weighting is never given as a number. Pathway B activates all four tests with no disclosed split, and B is likely the modal pathway – so most baseline scores will rest on an undisclosed weighting between a quantitative estimate and an analyst judgment.
- Pathway assignment partly tracks Sylvera's own data coverage. "High MSL coverage" is a property of your LiDAR acquisition footprint, so projects in under-flown geographies – much of Africa, and dry and open systems generally – are routed away from Pathway A irrespective of their own rigour.

Ask:

- Make the rule deterministic, so that every project lands in exactly one pathway. That means correcting the boundaries so a project at exactly SMD 0.4, or exactly 75% matched, is not in two pathways at once; stating whether both criteria must be met or only one; and stating what happens when the two disagree – good balance but poor match coverage, for example.
- Define what separates A from B. As written the dynamic baseline criteria are identical for the two, so the entire difference rests on "high MSL coverage" versus "coverage." Please define both terms and the threshold between them.
- Publish the weighting as numbers. For each pathway, state the weight given to the quantitative and to the qualitative assessment, and clarify whether the



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

greyed-out tests in A and C are excluded from the score entirely or simply down-weighted.

We would also refer back to section 2: in our view the gate itself should be predictive validation rather than covariate balance.

Sylvera Response:

The pathway logic is intended to be fully deterministic, and we apologise for any lack of clarity in the materials that suggested otherwise. A project remains in Pathway A only where both conditions are met: $SMD \leq 0.2$ (tightened from 0.4) and match coverage $\geq 75\%$. Failing either routes the project to pathway C.

What separates A from B is two distinct gates. A context gate captures circumstances our dynamic baseline structurally can't account for, such as high soil carbon or undisclosed land tenure changes. A quality gate reflects our confidence in Sylvera's underlying independent data: where our Biomass Atlas lacks Multi-Scale LiDAR coverage, for instance, we're less comfortable relying on the DB as the source of truth over carbon stock values. We recognise the disparity this can create for under-flown geographies like Africa and dry and open systems, and closing it is a priority for future Biomass Atlas updates. In the meantime, we aim for the same standard of assessment regardless of pathway, even if that means leaning more heavily on qualitative review or alternative data sources.

On the numeric weighting per pathway, treatment of greyed-out tests, and the section 2 point on predictive validation versus covariate balance, we'll set these out in the forthcoming whitepaper.

6. The donor pool has no eligibility criterion – only a geographic one. Protected areas are not excluded.

Donor pool and control eligibility

Comment: The donor pool is defined geographically – a ~100 km buffer intersected with ecoregion, less the leakage collar, other REDD+ projects, out-of-country administrative areas, non-forest and plantations. There is no eligibility criterion: no



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

requirement that controls share the project's tenure or governance regime, and no exclusion of protected areas or national parks, which operate under a categorically different enforcement regime and in which the project's intervention could never have been deployed. This is a treatment-assignment problem and it is directional. Controls drawn from land where deforestation is suppressed by protection rather than by absence of pressure depress the counterfactual, making the project's own baseline appear divergent on the high side and attracting an over-crediting flag. This is among the best-established threats to validity in quasi-experimental evaluation of conservation interventions (Ferraro & Pressey, 2015; Schleicher et al., 2019).

Ask: Add an explicit eligibility filter: controls must be land that could in principle have received the same intervention, on grounds of equivalent tenure, governance structure, ecological context and deforestation drivers. Exclude strictly protected areas, and publish the exclusion dataset and buffer used. Separately, we would welcome an empirical justification for the 5 km leakage collar.

Sylvera Response:

We're using the 5 km leakage collar in alignment with PACT v2.1 from the Cambridge Centre for Carbon Credits (4C) as a default, rather than deriving it independently for this pipeline. However, after assessing whether it's available, appropriate, and justifiable, our framework will usually use the project's own leakage boundary instead. In that same spirit of appropriateness testing, for each project we test a variety of donor zone sizes, including the project's own reference area, and select whichever gives the best match.

On donor eligibility, we've begun adding land tenure and governance structure datasets into the donor-zone masking across a number of jurisdictions (where available) and we now exclude all protected areas from the donor pool entirely, so controls drawn from land under a categorically different enforcement regime are no longer eligible. We recognise this is a limitation where the social-context match isn't perfect, and we've factored that into our pathways - we'll weight the dynamic baseline less where we and the project agree the match may be inappropriate.



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

7. Stopping active destruction along the Highest Deforestation Frontline is better than stopping invisible “threats or risks”, even though they are not yet visible on outdated Verra maps.

Put plainly: it is difficult to use the methodology to protect rainforest and stop active destruction, when the mandatory mapping system does not even show that very deforestation yet – and will only show it with a two-year delay. Verra’s forcing of outdated maps causes this problem – the fastest destruction is per se not visible on those maps, making extreme deforestation technically potentially inadmissible, and even in the best cases, subject to extreme crediting delays and uncertainties, and thus additional risk that those very projects should not need to shoulder. To enable those projects to combat the worst deforestation, rather than to hold them back, the Sylvera review framework (alongside the independent deforestation review) should continue to assess also recent deforestation and help consider data lag issues (including in carbon maps). In technical terms, the Biomass Atlas / Hansen inputs are coarse and lagged, and the framework itself leaves the most recent years “unconfirmed” – precisely when clearing on an active frontier is fastest – so the observed-loss counterfactual understates reality. Set the counterfactual on fresher, ground-truthed frontier data. In fact, the counterfactual becomes factual with current data of actual ongoing destruction. The assessment of “active destruction” then becomes replaced with “stopping active destruction”. This is what REDD should be all about, and Sylvera can continue to support this journey with a flexible “gap filling” framework, to help developers working in areas of extreme deforestation with data gaps, independent deforestation assessment, and flexibility on data gaps. As an example, RDC uses certified deforestation data that is up to date to create “gap analysis” (dated, certified deforestation alerts) to write recent clearing down before crediting and would provide it for validation.

Sylvera Response:

Our AUD ratings assess vintages on an ex-post basis: for each credited year, actual observed deforestation is compared against the dynamic baseline’s counterfactual for that same year. We recognise this still involves a reporting lag, and that even a dynamic, annually-refreshed baseline is not a real-time measure, but because both the project’s deforestation and the matched control zone’s deforestation are measured on the same annual cycle, the lag applies equally to both sides of the comparison.



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

Recent destruction isn't rendered inadmissible by this; it's assessed and credited once that monitoring year closes, on a like-for-like basis. It's also worth distinguishing our approach from jurisdictional risk-map methods such as VM0048, which model forward-looking deforestation risk: our dynamic baseline isn't a predictive risk measure at all, it tracks actual, observed annual background deforestation in comparable forest each year, which is closer to the "active destruction" framing, just measured empirically rather than via real-time alerts. Where gap-filling, near-real-time alert data is available, we will use it as corroborating evidence to support the assessment, but not as a substitute for the verified baseline, as independent verifiability is a key requirement of our REDD ratings. Any future increase in deforestation will be captured within the relevant vintage as and when it occurs and credited accordingly; we do not attempt to forecast this directly, given the well-documented volatility and difficulty associated with deforestation projections. That said, the risk of increased future deforestation continues to be assessed elsewhere within the framework, including within additionality, where relevant evidence is weighted accordingly.

8. Make prevented deforestation visible by design.

While the focus on stopping extreme, active destruction is primary, the focus should also include an ability for the project developer to showcase that they have prevented a worsening of the situation "that never made it into the maps" because it was successfully prevented. In the current form, under the presented framework, if a project succeeds, the avoided clearing never appears on any map. The map-only, observed-loss framework rewards only the projects that fail to stop destruction, and penalizes those that succeed. To overcome this, we recommend admitting forward-looking, documented threat evidence, such as road paving schedules, logging-concessions and land-regularization plans, and enforcement records, into the baseline, alongside the observed series, so that preventing these plans can be seen as a crediting event. In our case, for example, we are fighting plans to privatize a 35,000 ha area for logging.

Sylvera Response:

Our dynamic baseline is designed to make avoided deforestation visible. We don't need to observe the averted clearing inside the project itself; we infer it from the divergence



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

between the project's actual loss and observed loss in a matched control zone of comparable unprotected forest surrounding the project area. The more deforestation a project prevents relative to that control, the larger the divergence, and the more it's credited, so success isn't penalised, it's precisely what generates the credit.

On the 35,000 ha logging privatisation example, that scenario, with clear intent to harvest the project area under a business-as-usual case, would fall under Avoided Planned Deforestation (APD) pathway rather than AUD. For APD projects we do not apply the dynamic baseline test but rather look at the likelihood and scale of clearances in, considering if the economic activity would have been legal and is comparable to historical economic activities in the region based on proxy concession areas. Lastly, any deforestation drivers or context of the dynamic baseline misses (where we consider it well justified) would be reflected in the qualitative assessment. We see the dynamic baseline as an important tool to guard against overcrediting risk, but we recognise the nuance and other factors that go into REDD, where qualitative assessment of deforestation drivers also needs to be considered.

9. Let evidenced imminent drivers set the baseline – and, particularly in intact forest frontlines, treat legalized clearing as a problem, not protection.

A counterfactual matched only to recent loss ignores dated evidence of an upcoming surge, as in the case of the BR-319 paving, and State plans to grant logging rights and to regularize public forest (i.e., to convert illegally occupied public forest into titled private land, perhaps even in violation of legal requirements). Such plans raise both the baseline and additionality – the project is racing against the destructive plan – and should be counted as forest protection. On our frontier the dominant driver is illegal clearing (grilagem). Evidenced illegal, or soon-to-be-legalized, clearing should support baseline plausibility, not weaken it – as the framework would potentially do in its current form. Note also that under VM0048 our deforestation rate is Verra-allocated, not chosen – so any divergence from your dynamic baseline should be interpreted against that jurisdictional allocation, not scored as a developer decision.

Sylvera Response:



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

We assess threats and drivers on an ex-post basis, for each year credited, so a documented future plan alone (e.g a paving schedule, a concession announcement) would not raise the baseline ahead of any observed clearing. Where illegal clearing is already underway in the surrounding landscape, our dynamic baseline is designed to capture it as it happens: "distance to recent deforestation" and "distance to roads" are matching covariates specifically intended to reflect pressure on an active frontier, so that clearing is reflected in the counterfactual once it is observable, not only once it reaches the project boundary. Legalisation or regularisation of that clearing doesn't change this treatment, its legal status isn't the determining factor. What matters is location and timing: clearing targeting the project area itself, whether legal or illegal; clearing outside the project area is captured by the dynamic baseline once that threat actually materialises.

The avoidance of legal deforestation is assessed under the APD, and the avoidance of illegal or irregular deforestation under the AUD pathway.

With regards to VM0048, we recognise that a Verra-allocated rate is not a developer's choice, and understand the position this places developers in; any divergence would be attributed accordingly rather than read as a developer decision. We would, however, still flag a material overcrediting risk in commentary where it exists, consistent with how we treat any external risk.

10. Update road-distance data to reflect real, current road expansion.

Sylvera's matching uses historical road-network data to identify "similar" control plots. In corridors where new roads are being built or paved right now – like BR-319 adjacent to this project – that historical data is out of date. Parcels that look like "deep interior forest" in the legacy data are in fact frontier parcels facing active road-driven deforestation pressure. Matching them to true interior-forest controls understates the threat and under-credits the project. We recommend Sylvera flag projects in active road-expansion corridors and match on current frontier exposure, not legacy road distance.

Sylvera Response:

Road distance is derived from Overture, so recent construction is captured in principle. The



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

caveat we would add is that Overture's road geometry is largely OpenStreetMap-derived, and OSM is updated through contributor activity and periodic bulk imports rather than a uniform global resurvey. Coverage is therefore heterogeneous in space and lumpy in time. We will regularly update datasets used in our analysis to ensure the most relevant data is used.

ADDITIONALITY

Financial Additionality

11. Credit the at-risk cost of securing the forest – and separate genuine developers seeking real protection and impact where carbon finance is a critical tool from carbon project developers that are there for the revenue only instead.

On frontier rainforest areas facing extreme active destruction, the developer must defend and secure the forest before carbon rights can be finalized, facing a scenario of larger up-front burden and, when successful, a larger reward, with more forest standing when rights close. The framework should reward frontline investment and should not penalize in-progress carbon rights where the delay reflects market and institutional failures betraying the duty to protect – more so where institutional weakness effectively lets criminals take public forest – while still distinguishing serious developers with a real presence on the ground from those without one.

Sylvera Response:

The Sylvera framework does not automatically penalise upfront or overlapping finance, recognising the role it plays and the practical difficulty of initiating a carbon project ahead of credit issuance. Where overlapping finance is present, we ask the developer to provide evidence that carbon finance was material to enabling the continuation of protection activities.

On carbon rights, we assess each project according to its own sociopolitical context. However, effective and timely free, prior, and informed consent (FPIC) of Indigenous People's and local communities, and agreements with existing landowners, where relevant, to secure carbon rights ahead of project activities, is a widely recognised market standard, and one we



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

consider important for both effective safeguarding and long-term permanence built on genuine social buy-in. As with other elements of our framework, any risk we identify is based on our initial research and analysis, and developers have the opportunity to provide additional project context through the standard engagement process, which we take into account in finalising the assessment.

PERMANENCE

12. The Biomass Atlas is treated as insufficient for baseline work in some contexts but still relied on quantitatively in Permanence.

Biomass Atlas: accuracy, and consistency with Permanence

Our most concrete consistency concern: the framework accepts that the Atlas is not a sufficient basis for quantitative baseline work in certain ecoregions, that is what triggers the qualitative pathway – yet Permanence relies on the same data quantitatively, both to confirm reported carbon stock loss and to identify non-reported loss and reversals. If the Atlas is an unreliable proxy for biomass change in a given forest type, it should not simultaneously be treated as sufficient to characterise reversals there.

Ask: Publish per-biome confusion matrices with sample sizes derived from the Cochran equation or equivalent, reporting overall, user and producer accuracy, with dry and open woodland broken out – the regime where 30 m AGB-derived forest/non-forest classification is weakest. Where the Atlas is judged insufficient for quantitative baseline assessment in a region or forest type, apply that judgment consistently across Permanence and set out how reversals are assessed there. Please also confirm that the forest definition used in the pipeline matches the project's registered or national definition, and that the identical definition, minimum mapping unit and dataset are applied to project and control plots, so classification error introduces no directional bias between them.

Sylvera Response:

Where BiomassAtlas is deemed unsuitable for a region under REDD V.4, it is excluded from the permanence score for that region. Where BiomassAtlas is not used, we still evaluate other data sources judged appropriate to the context to ensure any genuine permanence



Website sylvera.com



Email
projectratings@sylvera.io

7 September 2026

risk is captured. For fire, this can include national monitoring systems where available, third-party datasets such as MapBiomass Fire, and the programme's own reporting, with each source assessed on its appropriateness and quality before being relied upon.

Regarding the forest definition, we will configure our definition to the national definition, with any mismatches flagged appropriately.

SAFEGUARDING AND CO-BENEFITS

No Comments

OTHER COMMENTS

13. Risk Differentiation (Project vs. Jurisdictional): How will the framework accurately differentiate between project-specific risks and systemic risks originating from jurisdictional or national REDD+ programs, which remain strictly outside the project's direct control?

Sylvera Response:

Any risk to delivery, safeguarding, permanence, or additionality is treated as a risk to the project regardless of whether it originates internally or externally, since materiality to the project's outcomes is what determines relevance to the rating, not the source of the risk. Where a risk is judged material, it is flagged in the assessment. Developers are engaged throughout the process to ensure external nuance is properly understood and reflected. Where a material risk is included, the commentary should clearly attribute it as either project-driven or external in origin, so the distinction is transparent to the reader. If a developer considers this attribution inaccurate or unfairly worded, it can be revisited and amended through the standard developer engagement and review process.



Website sylvera.com



Email
projectratings@sylvera.io