ARBOREAL CAPITAL

Research Note 005
The Verification Problem

How to confirm that a tree is alive at a specific coordinate, at global scale, for centuries, without a central authority. Satellite infrastructure, ground truth networks, the alive-until-proven-dead default, and what happens when the satellites stop working.

Mitch Hart

March 2026

---

1. The Core Requirement

The Arboreal Capital mechanism ties token value to tree survival. This creates a conservation incentive only if the system can credibly determine whether a given tree is alive. A token holder must believe that if the tree dies, the system will discover it and freeze the token. Without that belief, the token is a fiction—a number attached to coordinates with no relationship to physical reality.

Verification must satisfy five constraints simultaneously. It must operate at global scale across billions of trees distributed over every continent except Antarctica. It must function without a central authority, because any central authority is a single point of failure with a finite lifespan. It must be accurate enough that the market trusts it, but it need not be perfect—the mechanism tolerates some error, as explained below. It must be affordable, funded by the system itself rather than by external grants or institutional budgets. And it must persist across centuries, surviving the disappearance of any individual technology provider.

> **Design Constraint:**
>
> *The verification system must confirm tree survival at global scale, without central authority, at a cost the mechanism can self-fund, for longer than any existing institution has survived.*

2. The Three-Layer Architecture

No single verification method satisfies all five constraints. Satellite imagery is global but imprecise. Human inspection is precise but expensive and unscalable. Sensor networks are accurate but require physical infrastructure that degrades and must be maintained. The solution is a layered architecture where each method compensates for the weaknesses of the others.

2.1 Layer 1: Satellite and Aerial Imagery

The base layer uses satellite and aerial imagery to monitor canopy presence at registered coordinates. Multiple commercial providers—Planet Labs, Maxar, Airbus Defence and Space, and others—offer multispectral imagery at resolutions ranging from 30 centimeters to 3 meters, with revisit rates between daily and weekly depending on latitude and cloud cover.

For bulk trees in dense forest, satellite verification answers a coarse question: is there still canopy at these coordinates? Normalized Difference Vegetation Index (NDVI) analysis can detect canopy loss, distinguishing live vegetation from bare ground, dead wood, or structures. This is sufficient for the vast majority of registered trees. A boreal spruce in Siberia does not need individual identification; it needs confirmation that the forest at its coordinates has not been cleared.

For hero trees—named specimens of exceptional age, size, or cultural significance—satellite imagery serves as a first screen, not a final determination. A hero tree with a $500,000 market cap warrants higher verification fidelity than a bulk tree worth $2. The cost of verification should scale with the value at stake.

2.2 Layer 2: Ground Truth Network

The second layer is human verification. Token holders, hikers, park rangers, ecologists, and anyone with a smartphone and GPS can submit ground truth reports for trees near their location. A ground truth report consists of a geotagged photograph, a timestamp, and the reporter’s staked attestation that the tree at the given coordinates is alive.

Ground truth does not replace satellite imagery; it calibrates it. A satellite pass that flags potential canopy loss at a coordinate triggers a ground truth request. If a nearby holder or community member visits the site and confirms the tree is alive, the flag is resolved. If the ground truth report confirms the loss, the death claim process described in Research Note 002 is initiated.

The ground truth network is inherently biased toward accessible trees in populated areas. Trees near cities, trails, and roads will receive frequent ground truth verification. Trees in roadless Siberian taiga will receive almost none. This is acceptable for two reasons. First, accessible trees face higher anthropogenic threat—logging, development, vandalism—and therefore benefit most from frequent verification. Second, remote trees in undisturbed wilderness are overwhelmingly more likely to be alive than to have been secretly removed. The probability distribution is asymmetric, and the verification architecture should reflect that asymmetry.

2.3 Layer 3: Statistical Spot-Checking

The third layer is random audit. Each verification cycle, a small percentage of registered trees are selected at random for intensive verification regardless of whether any anomaly has been flagged. These spot-checks may involve high-resolution commercial satellite tasking, drone overflights, or commissioned ground inspections depending on the value and accessibility of the tree.

The statistical logic is borrowed from quality control in manufacturing. You do not inspect every unit coming off an assembly line. You inspect a random sample at a frequency that gives you a known confidence level. If the sample passes, you have statistical confidence in the batch. If the sample fails, you increase inspection frequency for that region or tree class.

The key insight is that the spot-check rate need not be high. If a bad actor knows that any registered tree has a 2% chance of being audited in any given year, and that a failed audit triggers cascading investigation of nearby trees, the expected value of fraud collapses. The deterrent is the uncertainty, not the frequency.

Key Insight:

Verification does not require certainty. It requires that the probability of detection is high enough to make fraud unprofitable. A 2% annual audit rate with cascading investigation achieves this.

3. Alive Until Proven Dead

The system’s default assumption is that a registered tree is alive. This is not optimism; it is the correct Bayesian prior. Old-growth trees that have survived 250 years have demonstrated extraordinary resilience. The annual mortality rate for old-growth trees in undisturbed forest is estimated at 0.5–1.5% depending on species and region (Stephenson et al., 2011). In any given year, the overwhelming probability for any individual tree is that it is still alive.

This default matters because it determines the burden of proof. A death claim requires affirmative evidence. A lack of recent verification does not trigger a freeze. If satellite coverage is temporarily unavailable due to cloud cover, sensor failure, or provider disruption, the token continues to trade normally. The tree is presumed alive until someone demonstrates otherwise.

The alternative—requiring periodic proof of life—would create a catastrophic failure mode. If verification infrastructure goes offline for six months, every tree in the system would be flagged as unverified. Token prices would collapse not because trees died but because a satellite company had a billing dispute. The alive-until-proven-dead default makes the mechanism robust to temporary verification gaps.

There is a legitimate concern that this default allows a dead tree to remain tradeable for some period before detection. This is true, and it is acceptable. The purpose of the mechanism is not to provide real-time mortality data. It is to create a durable economic incentive for conservation. A tree that dies and is detected six months later still served its conservation function for every year it was alive and tokenized. The token holders who bought during the undetected period bear the loss, which is a market risk disclosed in the token’s design—no different from buying stock in a company that has already filed for bankruptcy but has not yet announced it.

Design Principle:

A tree is alive until proven dead. Absence of verification is not evidence of death. The system tolerates detection lag because the conservation incentive operates continuously, not at the moment of verification.

4. The Provider Independence Problem

The verification architecture described above depends on commercial satellite imagery providers. As of 2026, the commercial Earth observation market is dominated by a small number of companies: Planet Labs (approximately 200 satellites), Maxar Technologies, Airbus Defence and Space, and a growing constellation of smaller operators. The combined infrastructure is robust, but any individual provider can go bankrupt, be acquired, change pricing, or restrict access.

A conservation mechanism designed for centuries cannot depend on the continued existence of any commercial entity. The mitigation is structural redundancy at every level.

First, the protocol specifies verification standards, not verification providers. The smart contract does not reference Planet Labs or Maxar by name. It accepts a death claim supported by evidence that meets a defined standard—multispectral imagery at a minimum resolution, ground truth photographs with GPS metadata, or independent arbiter confirmation per Research Note 004. When a provider disappears, another provider that meets the standard takes its place. The standard persists even if every current provider does not.

Second, the ground truth network provides a fallback that requires no commercial infrastructure at all. A human being with a camera can verify a tree. This technology—walking to a location and looking at it—has been available for approximately 300,000 years and is unlikely to become unavailable. If every satellite in orbit deorbits tomorrow, the ground truth layer still functions.

Third, the alive-until-proven-dead default means that a gap in satellite coverage does not crash the system. If commercial imagery becomes unavailable for a period—due to market consolidation, geopolitical conflict, or Kessler syndrome rendering low Earth orbit unusable—the mechanism continues to operate. Verification quality degrades. Detection lag increases. But tokens continue to trade, the conservation incentive remains operative, and the system waits for the next generation of observation technology to emerge.

Design Constraint:

The verification system must survive the bankruptcy of every current satellite operator. It does so by specifying standards rather than providers, maintaining a zero-infrastructure human fallback, and defaulting to alive during coverage gaps.

5. Verification Economics

Verification costs money. The question is whether the mechanism generates enough revenue to fund its own verification indefinitely. This paper addresses whether the architecture works; Research Note 006 addresses whether the economics work. But the two questions interact, because the cost of verification determines the minimum viable market capitalization of the system.

At scale, satellite imagery costs fall dramatically. Planet Labs offers monitoring at approximately $1–2 per square kilometer per year for programmatic access. A tree occupies roughly 100 square meters. Even assuming dedicated tasking rather than archive imagery, the marginal cost of satellite verification per tree approaches zero at sufficient scale. The binding cost is not imagery acquisition; it is imagery analysis—the computational and human labor required to determine whether a change in canopy at a registered coordinate represents a dead tree, a seasonal change, cloud shadow, or sensor artifact.

This is a machine learning problem that is already substantially solved for deforestation monitoring. Global Forest Watch, operated by the World Resources Institute, processes Landsat and Sentinel imagery to detect tree cover loss at 30-meter resolution globally, funded by institutional grants. The Arboreal Capital mechanism does not need to build this capability from scratch. It needs to fund access to it—or to its successors—in perpetuity.

Ground truth verification has a different cost structure. The marginal cost is the value of a person’s time to visit a location and take a photograph. For trees near population centers, this cost is effectively zero—people visit forests for recreation and can verify trees as a byproduct. For remote trees, the cost is significant: expedition logistics, travel, and time. The spot-check budget must account for this by reserving a higher per-tree allocation for remote and high-value specimens.

6. Edge Cases

6.1 Slow Death

Trees do not always die suddenly. A tree can lose branches over decades, suffer progressive disease, or be slowly killed by changes in hydrology. At what point is it dead? The protocol defines death as the irreversible cessation of biological function in the primary stem. A tree that has lost 90% of its canopy but retains living cambium is alive. A standing dead trunk is dead. Ambiguous cases are resolved through the arbitration process described in Research Note 004.

6.2 Coordinate Drift

GPS coordinates are not perfectly stable over geological time. Tectonic movement shifts coordinates by centimeters per year. Over 500 years, a tree in the Pacific Ring of Fire could shift several meters relative to its registered coordinates. The exclusion radius around each registered tree (specified in Research Note 003) provides a buffer. As long as the tree remains within the exclusion radius of its registered coordinates, verification can confirm its presence. If tectonic drift or geodetic datum revisions shift coordinates beyond the buffer, a coordinate correction dispute (Research Note 004) can update the registration.

6.3 Verification in Conflict Zones

Trees in active conflict zones may be inaccessible to ground truth verifiers and may have restricted satellite coverage. The alive-until-proven-dead default handles this: the token continues to trade at whatever price the market assigns given the uncertainty. If a war destroys a forest, the trees are dead and the tokens will eventually freeze when evidence becomes available. The lag between destruction and detection is a market risk, not a system failure. Token holders in conflict-zone trees are implicitly taking a geopolitical risk, which should be reflected in lower token prices—a form of risk pricing that the market handles without protocol intervention.

6.4 Climate Change and Biome Shifts

Climate change is altering the geographic ranges of tree species. A tree that was well-adapted to its location 250 years ago may now be at the edge of its viable range. The verification system does not need to predict this. It simply observes whether the tree is alive. If climate change kills the tree, the token freezes. If climate change stresses the tree but it survives, the token continues to trade. The market may price in climate risk for trees in vulnerable biomes, creating a discount that reflects the probability of climate-driven mortality—again, a form of risk pricing that requires no protocol intervention.

7. Conclusion

The verification problem is often presented as the fatal weakness of any mechanism that ties digital assets to physical reality. How can you possibly monitor billions of trees across the entire planet? The answer is that you do not need to monitor all of them all of the time. You need a layered system where satellite coverage provides broad monitoring, human ground truth provides calibration and precision, and random spot-checks provide statistical confidence. You need the correct default—alive until proven dead—so that gaps in coverage degrade gracefully rather than catastrophically. And you need provider independence so that the system survives the disappearance of any individual technology company.

The verification system does not need to be perfect. It needs to be credible enough that the market trusts it, robust enough that it degrades gracefully, and cheap enough that the mechanism can fund it in perpetuity. The three-layer architecture, the alive-until-proven-dead default, and the standards-based provider independence described in this paper are designed to satisfy all three requirements.

Key Insight:

The verification system does not monitor every tree. It creates conditions under which the probability of detecting a dead tree is high enough, and the cost of fraud is large enough, that the market treats tokens as credible claims on living trees. Credibility, not omniscience, is the design target.

References

Crowther, T.W. et al. 2015. Mapping tree density at a global scale. *Nature* 525, 201–205.

Hansen, M.C. et al. 2013. High-Resolution Global Maps of 21st-Century Forest Cover Change. *Science* 342(6160), 850–853.

Planet Labs. 2024. Planet Monitoring: Technical Specifications. Planet Labs PBC.

Stephenson, N.L. et al. 2011. Causes and implications of the correlation between forest productivity and tree mortality rates. *Ecological Monographs* 81(4), 527–555.

World Resources Institute. 2024. Global Forest Watch: Technical Documentation. World Resources Institute, Washington, DC.

Arboreal Capital. 2026. Research Note 002: The Death Profit Paradox.

Arboreal Capital. 2026. Research Note 003: The Gold Rush.

Arboreal Capital. 2026. Research Note 004: The Arbitration Problem.

Arboreal Capital · arborealcapital.io · Research Note 005 · v1.0 · March 2026