



Agentic Artificial Intelligence and its implementation in Transforming Audit, Policy and Environmental Accounting

Anirudh Khajuria

Indian Institute of Forest Management, Bhopal, India.

To Cite this Article: Anirudh Khajuria, "Agentic Artificial Intelligence and its implementation in Transforming Audit, Policy and Environmental Accounting", Indian Journal of Computer Science and Technology, Volume 05, Issue 01 (January-April 2026), PP: 326-333.



Copyright: ©2026 This is an open access journal, and articles are distributed under the terms of the [Creative Commons Attribution License](#); Which Permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract: Agentic AI is a monumental development in enterprise computing, moving away from traditional role-driven query systems to independent autonomous, goal-oriented actors who make decisions, plan and learn continuously. This article oceans the implications of agentic AI on four fields where there are stakes involved: audits and compliance; policies and regulatory automation; environmental accounting and GHG measurement systems; and forestry carbon economic systems. With existing market analyses, this article will provide a substantive, multi-domain analysis on agentic AI current capabilities and level of economic impact and what implementation challenges may arise as well as the future direction that will culminate with contractor-generated market revenue of \$USD 7.06 Billion in Year 2025 to \$USD 93.2 Billion by Year 2032 at an annual compound growth rate of roughly 43%. Agentic AI will therefore bring about a complete transformation within the professional services marketplace with respect to both environmental governance methodologies and the structure of carbon markets. of carbon markets.

Key Words: Agentic AI, Autonomous Agents, Environmental Accounting, GHG Measurement, Carbon Credits, Audit Automation, Policy Compliance, Forestry Carbon, Consulting Transformation

I. INTRODUCTION

The Paradigm Shift: From Reactive Tools to Autonomous Agents

Traditional artificial intelligence (AI) systems work as advanced input-output machines, taking input (the query), processing it through pre-trained models (their analytics), and then producing an output back to the user of the system (the AI's response). These types of systems are technically impressive, but in general, they are reactive. They do not begin actions (initiate actions) to achieve goals over time, nor can they adjust their methods (strategies) based on information about the environment (feedback from the environment) while pursuing a defined target.

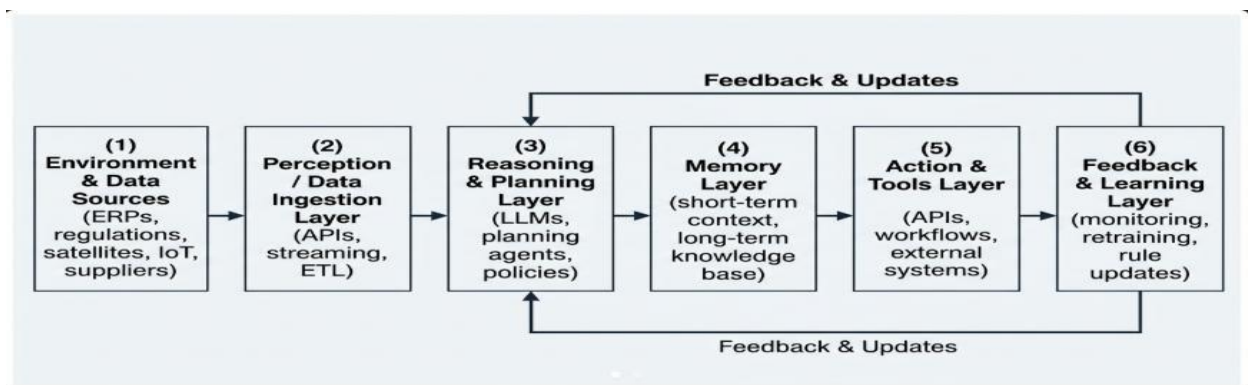


Figure 1. High-level agentic AI architecture showing the perceive-reason-act-learn loop across enterprise systems. (geeks for greeks)

Agentic AI breaks the mold of traditional AI implementations entirely. There are four defining features of agentic AI systems.

- **Permanence:** Agentic AI systems can operate continuously and without interruption for long periods of time without needing a person to restart them.
- **Autonomy:** Agentic AI systems are capable of making independent decisions in accordance with defined goal structures without waiting for explicit instructions.
- **Proactivity:** Agentic AI systems have the ability to take initiative and meet needs before the need for that action arises, rather than waiting for instruction from another person.
- **Social Ability:** Agentic AI systems can communicate with people and other AI systems in natural language in a multi-agent environment.

The combination of these characteristics changes the way in which AI is viewed—from an analytical tool to an active participant in the organizational process. For example, a traditional analytical platform may detect a problem in financial data (e.g., by identifying anomalies), while an agentic AI system will investigate the cause of the anomaly, review the applicable regulations, draft a corrective action memo, and notify the appropriate compliance officer—none of which would have occurred had no person made the request.

1.2 Consulting – The Proving Ground

The consulting industry has an enormous opportunity to use agentic AI because the industry has an inherent contradiction: the main source of value from consulting firms (i.e., expert human judgement) is also the most expensive input. For example, a senior consulting partner who bills at \$500 - \$1,000 per hour will only provide value to clients for a limited number of hours each day. Increasingly complex regulations are compounding this issue by requiring ongoing monitoring that the size of teams cannot provide sustainably at scale (e.g., EU Corporate Sustainability Reporting Directive (CSRD), SEC Climate Disclosure Rules, and evolving GHG Accounting Standards).

Agentic AI can resolve this paradox by decoupling the expertise of humans from the time required of them, encoding the expert logic of those "human" consultants into persistent, scalable agents that can continually execute that logic. This creates an opportunity for those "human" consultants to free up time to focus on higher-value strategic work, client relationship-building, and complex judgement that machines cannot replicate.

1.3 Scope and Methodology

This document aggregates public views on seven categories: Including agentic systems architecture & design; auditing transformation; automated policy & regulatory compliance; GHG measurement & reporting; carbon forestry economics; implementation strategies and; future tech trajectories. The data sources consisted of primary market research from MarketsandMarkets, Mordor Intelligence, Precedence Research, GM Insights, Grand View Research, & Climate Focus’ Voluntary Carbon Market ’24 Review.

The agentic AI market has recently started experiencing hypergrowth. According to MarketsandMarkets, it was valued at USD 7.06B in 2025 and is predicted to grow to USD 93.20B by the end of 2032, representing a compound annual growth rate (CAGR) of 44.61%. By contrast, Mordor Intelligence provides an estimate of USD 57.42B for 2031 and a CAGR of 42.14%. Precedence Research projects an upper limit of USD 199.05B by 2034 and a CAGR of 43.84%. the amelioration of generative AI, reinforcement learning and orchestration framework such as langchain, crewai and Autogen has created infrastructure upon which agentic applications can be build and deployed .the agetic market was valued at USD 5.43 billion in 2024 and is forecasted to reach USD 236.03 BILLION at CAGR of 45.82%.

Source	2025 Value	Forecast Year	Forecast Value	CAGR
MarketsandMarkets	USD 7.06 bn	2032	USD 93.2 bn	44.6%
Mordor Intelligence	USD 9.89 bn	2031	USD 57.4 bn	42.1%
Precedence Research	USD 10.86 bn	2034	USD 199.1 bn	43.8%
Fortune Business Insights	USD 7.29 bn	2034	USD 139.2 bn	40.5%
Dimension Market Research	USD 7.4 bn	2034	USD 171.2 bn	41.7%

1.4 Sectoral Distribution of Demand of Agentic Work

All sectors don't have equal levels of demand/deployment. The banking/financial services/insurance (BFSI) industry accounts for the majority of agentic AI currently deployed because of their need for fraud detection, regulatory compliance and automated underwriting. The environmental/sustainability (ESG) market sector is also on track to become one of the fastest-growing verticals due to regulatory changes that require mandatory reporting on ESG criteria, and also because of the rapid growth of the voluntary carbon markets. The other key markets are made up of healthcare, manufacturing and IT or telecoms.

The product split of software components dominates the agentic AI marketplace to 50.4% of total market share in 2025, with the remaining balance made up of professional services / implementation support.

II. AUDIT TRANSFORMATION THROUGH AGENTIC AI

2.1 Structural Limitations of Traditional Audit

Traditional audit methodologies have three main interconnected limitations. Firstly, there is sampling bias, as auditors only examine a small percentage of total transactions (typically between 5 to 15 percent), leaving a much larger group of transactions (the majority) unaudited and leaving a strong possibility that there are material errors. Additionally, there is temporal lag through annual or quarterly audits, where gaps exist (lots of time) prior to when the audit is performed, allowing for fraud/errors/non-compliance to occur over time before they are caught. Finally, there is a scalability ceiling, where deployable audit staff resources cannot scale at a normal rate, which is required to perform audits, and rightly so, as data volumes per entity, per system, are now approaching terabytes on a daily basis.

In addition to creating operational challenges for organizations, these audit limitations create substantial risk in an organization. For example, the Association of Certified Fraud Examiners (ACFE) estimates that organizations worldwide lose approximately 5 percent of their annual revenues due to fraud and that a major reason for this is due to failing to detect fraud in real-time through periodic audits.

2.2 The Agentic Audit Architecture

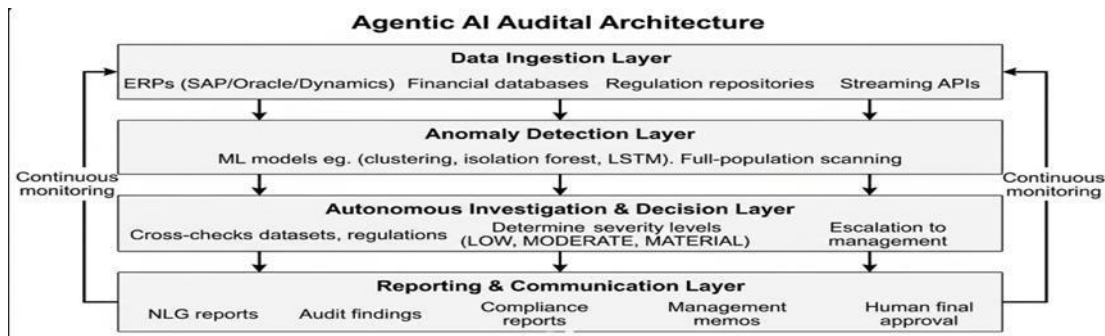


Figure 2. Agentic AI audit architecture with continuous data ingestion, ML-based anomaly detection, autonomous investigation, and automated reporting. (see the generated image above (ISACA))

Agentic AI audit solutions use technical architecture with 4 distinct functional components:

1. There is a data ingestion layer which provides persistent API connections to ERPs (SAP, Oracle, Dynamics), financial databases, regulation repositories and other external data and provides continuous streams of data instead of only providing periodic data exports.
2. The anomaly detection layer scans the transaction population through machine learning trained models (i.e., unsupervised clustering algorithms, isolation forests and LSTM neural networks) to identify abnormal transaction activity as compared to what had been expected. Machine Learning trained models are more effective against evolving fraud schemes and changes in behaviour within the organisation when compared to only using a rule-based system approach.
3. If an anomaly is detected, the agentic AI will begin autonomously investigating additional sources of data using a number of different compliance checks, including querying other datasets, confirming with regulatory requirements, etc. Any anomaly that is not considered material will be resolved through a LOW to MODERATE severity level, and any anomalies that are considered material will be escalated with pre-determined documentation to an appropriate final approval by management instead of authoring the final approval.
4. The reporting and communication layer uses natural language generation of content to produce audit findings, compliance reports and management communication automatically; therefore, the only aspect of final approval is a human approving it, and there will be no human authoring the documentation.

2.3 Economic Impact Analysis

There is a strong economic case supporting the notion of transforming the audit through agentic audits. For example, for a Fortune 500 company, the annual cost of providing external audit services is usually \$10-50 million per year when you consider the cost of providing external audit services, internal audit staff, audit technology infrastructure costs (e.g. computer hardware, software, etc.), and the time spent by management to provide external audit services. Agentic systems represent a significant upfront financial investment (usually between \$2-10 million for implementation, depending on the size and complexity of your organization), and they will significantly reduce your ongoing operational costs to approximately 10-20 percent of what you would have spent to conduct traditional audits. The impact of this cost inversion will be profound on a number of audit firms' business models. By developing agentic AI, an audit firm will be able to provide continuous monitoring services at a price point that is competitive with traditional periodic audits.

Additionally, the anticipated benefits of using agentic AI—forewarning of issues earlier than traditional audits, comprehensive coverage instead of 10 percent coverage of transaction populations (e.g. financial transactions), and the provision of real-time visibility of risks in the financial systems—will create a competitive differentiator that will justify charging more for audit services than what is charged for conducting traditional audits.

According to TCS research, agentic AI will transform all aspects of the audit value chain, including business development (automated prospect identification through financial statement analysis), engagement execution (AI-powered risk assessments, automated evidence ingestion and ML-based anomaly detection), and continuation of client relations (client benchmarking and regulatory compliance monitoring).

III. POLICY AND REGULATORY AUTOMATION

3.1 The Complexity Challenge

For the last ten years there has been a fundamental shift in the global regulatory environment. For example, the EU's new CSRD requires over 50,000 companies in its 27 Member States to publish detailed sustainability reports. Another example of the change is the SEC's rule that publicly traded US companies are responsible for reporting scope 1, scope 2 and scope 3 emissions. A third example of new global regulations are the regulators from over 30 different jurisdictions that now reference the TCFD framework for climate-related financial disclosures. The combined burden of compliance impacts multiple regulations from different jurisdictions and has thousands of overlaps all resulting in inconsistent requirements that are difficult for multinational enterprises to satisfy. It is impossible for a human compliance team to maintain sufficient regulatory

competence across this complexity and to do so at an acceptable cost. A large consulting firm could employ hundreds of specialists, but they specialize in limited jurisdictions and topic areas and are still likely to have gaps in their coverage as regulations change faster than they can be trained to comply with.

3.2 Dynamic Compliance Through NLP

Three Ways Agentic AI Overcomes This Challenge.

- 1) Automatic Regulation Extraction** Using Natural Language Processing: Advanced Language Models (consistent with Transformer Architectures like BERT, GPT & other Domain-Based Models) can quickly parse regulatory text, identifying obligations, prohibitions, thresholds and filing/reporting requirements with a high degree of confidence. When there is a new regulation that has been issued, agentic systems will parse, interpret and map the regulation's requirements to the agency's organizational structure in matter of hours as opposed to the many weeks required by current manual processes.
- 2) Automated Regulation Mapping:** Once extracted, each regulatory obligation will automatically be mapped to its applicable business processes, data systems and organizational units. This provides an ongoing compliance matrix that is continuously updated as regulations change, superseding the current static spreadsheet compliance trackers that become obsolete a few days after their publication once new regulations are issued.
- 3) Predictive Violation Risk Scoring:** Through machine learning algorithms, historical compliance patterns, current organizational activities, and corresponding regulatory requirements are analyzed to generate predictive violation risk scores based upon the likelihood of an organization violating a given regulatory obligation. Areas identified with high risk will automatically escalate to compliance teams' attention so the team can focus on mitigating only those items that present a genuine risk and not waste time performing periodic reviews that offer no true value.

3.3 Cross-Jurisdictional Harmonization

Multinational companies can potentially gain the most benefit from agentic AI by using it to align compliance with regulations across multiple jurisdictions. For example, when the EU brings in new ESG mandatory disclosure requirements which are partially overlapped by existing SEC rules, there will need to be a human compliance team that will reconcile the requirements, to identify gaps and make changes to the compliance procedures of all affected entities. However, the agentic AI system will perform this reconciliation automatically and will identify conflicts (defined as instances in which one jurisdiction has a requirement that conflicts with another) and opportunities (defined as instances in which compliance with one jurisdiction's requirements will also result in compliance with another jurisdiction's requirements), thereby significantly reducing the Compliance Overhead across the enterprise.

IV. ENVIRONMENTAL ACCOUNTING AND GHG MEASUREMENT

4.1 The GHG Measurement Challenge

Because greenhouse gas accounting presents such distinct technical issues, it is a very good fit for intervention by agentic AI. GHG emissions differ from financial accounting, which is based on discrete, exact transactions that are recorded in an organization's enterprise systems.

GHG emissions, in contrast, are spread across many areas in both physical and virtual locations (i.e., digital) and are continuously occurring; because of this, GHG emissions are often estimated and may represent multiple scales of activity across organizations, geographically, and temporally.

Using the GHG Protocol Scope framework as a basis for measurement in multiple ways introduces many challenges in measuring and accounting for GHG emissions. For example:

- **To measure Scope 1** emissions (which are the direct emissions that come from the GHG source itself), data regarding processes that create emissions across plant operations, vehicles, and fugitive emissions will all be required.

- **To measure Scope 2** emissions (the indirect emissions that come from the use of electricity and/or fossil fuels), data will be required to identify the energy used to create electricity and/or fossil fuel, depending upon the source.

To measure Scope 3 emissions (which are the indirect emissions impacting an organization that occur as a result of activities that are not within the organization's direct control), organizations will need to rely upon supplier engagement data. However, this type of data is rarely standardized, so many organizations will have to rely upon estimates using industry average emissions factors.

At present, traditional GHG accounting practices require organizations to create an annual inventory of GHG emissions by conducting a manual accounting exercise that is highly labor-intensive, and the results of which are often available 12–18 months after the end of the measurement period. Additionally, most traditional GHG accounting practices rely upon emission factors that may not accurately represent actual emission intensity for the GHG source in question.

4.2 Agentic AI Solutions for GHG Accounting

Automatically collecting data from multiple sources: Through constantly-connected agents to utility APIs, Fuel Purchase Records, Production Scheduling Systems, Transportation Management Systems, and supplier databases, this procedure will allow for the continuous collection of data, thus enabling the near-real-time tracking of process-level emissions instead of annual facility-wide emissions.

Applying Emission Factors Dynamically: Using machine learning models to keep continuously updated libraries of emission factors will allow for the automatic application of those factors to the activity data as the activities are being collected, thus removing the manual element of selecting emission factors, which creates lag times and inaccuracies with traditional emissions inventories. These emission factors will include grid emission factors for electricity; supplier-specific intensity factors based on

publicly disclosed supplier emissions data; and process-specific emission factors derived from the operational measurements collected from the energy supply process.

Estimating and Validating Scope 3 Emissions: Since measuring Scope 3 emissions in the supply chain is usually impossible, agentic AI has developed an accurate estimation method using sophisticated models to triangulate between supplier emissions data submitted by suppliers as required and verified by credible independent sources such as industry benchmarks, and empirical evidence from the physical supply processes using the same techniques as empirically measuring Scope 2 and energy supplied to customers. All estimate of Scope 3 emissions will include an appropriate level of uncertainty and will be prioritized based on the materiality of each supply chain contributor so that suppliers contributing to the largest amount of uncertainty can be engaged with the least amount of effort to facilitate the highest amount of accuracy.

Continuously Validating Emissions Inventories: Agentic systems will provide continuous external validation of emissions inventory data through internal checks and the use of external benchmarks and regulatory compliance checklists. Agentic systems continuously exceed the level of risk tolerance established by each of their owners and in doing so, provide early identification of data inaccuracies or errors before they are reported.

4.3 Regulatory Drivers

Agentic GHG accounting is also backed by several current regulatory developments. The EU's CSRD has created an obligation for companies to provide assurance over their sustainability information, which means businesses need to have a GHG data solution that is auditable and traceable (traditional systems based on spreadsheets don't provide for this). The SEC's climate disclosure ruling also compels companies to disclose material climate risks and, for very large companies, scope 1 and 2 emissions providing limited assurance as stated above; thus driving investments into automated, auditable GHG data infrastructure. What is more, automated agentic systems providing auditable continuously validated GHG data with complete calculation audit trail meet these requirements for an overall infrastructure solution.

V. FORESTRY CARBON AND CARBON CREDIT MARKETS

5.1 Carbon Credit Market Context

The voluntary carbon market is a form of financial system used to allow businesses to offset their carbon emissions in an effort to meet climate change goals. According to GM Insights, the overall size of the global VCM market was estimated at \$1.7 billion in 2024 and is expected to be worth \$15.7 billion by 2034, with a compound annual growth rate (CAGR) of 25 percent. Grand View Research has a more optimistic short-term estimation of the global VCM at \$4.04 billion in 2024, with a long-term forecast of \$23.99 billion in 2030, due to projected 35.1% CAGR growth rates. Roots Analysis provides even higher growth projections, with their forecast from \$1.6 billion in 2025 to \$47.5 billion by 2035 representing a CAGR of 38 percent from 2013 levels. Most of the VCM's total credit issued has been through forestry credits such as REDD+, afforesting and reforestation (or A/R) and improving the management practices on existing forests (or IFM) which account for approximately 40-50% of the total credits issued within the VCM in recent times.

Source	2024 Value	Forecast Year	Forecast Value	CAGR
GM Insights	USD 1.7 bn	2034	USD 15.7 bn	25%
Grand View Research	USD 4.04 bn	2030	USD 24.0 bn	35.1%
Roots Analysis	USD 1.6 bn (2025)	2035	USD 47.5 bn	38%

5.2 Technical Challenges in Forest Carbon Measurement

Carbon accounting for forests is one of the most difficult tasks for environmental science to perform. The challenges encompass biological, physical, and methodological aspects:

Forest carbon density varies on many different levels and factors including species, age structure, health, soil types, precipitation patterns, and management history. One management unit for a forest may include dozens of grey location types, each with different carbon dynamics and measurements.

It is much more difficult to measure below ground biomass and soil organic carbon compared to above ground biomass (which can be quantified using allometric equations based on diameter and height measurements). Below ground biomass and soil organic carbon represent 40 – 60% of the total carbon found in that forest.

To represent an actual climate benefit, a carbon credit must be able to demonstrate additionality (the carbon benefits produced in this project would not have occurred without it) and permanence (the sequestered carbon will remain sequestered). In order to confirm these two properties, the verification process must compare against a baseline for each of these properties and that baseline is not able to be directly observed (it represents what would have existed had that alternative forest scenario occurred). forest carbon—are considerably more difficult to measure and are subject to greater methodological uncertainty. **Additionality and Permanence:** For carbon credits to represent genuine climate benefit, projects must demonstrate additionality and permanence. Verifying these attributes requires counterfactual analysis against baselines that cannot be directly observed—they represent the forest that would have existed under an alternative scenario. **Measurement, Reporting, and Verification (MRV) Costs:** Traditional forest carbon MRV relies on field inventory teams that conduct plot-based measurements on 3–5 year cycles. For large REDD+

projects covering millions of hectares, MRV costs can reach USD 5–15 per hectare per year, consuming a substantial fraction of project revenues. The costs associated with Measurement, Reporting, and Verification (MRV) of carbon in forests have historically relied upon field inventory teams conducting plot-based measurements every 3 to 5 years. Consequently, MRV costs for large REDD+ projects covering millions of hectares can be anywhere from USD 5 to 15 per hectare annually, accounting for a large proportion of project revenue.

5.3 Agentic AI in Forest Carbon MRV

Agentic AI is changing the way forest carbon MRV will be accomplished through the use of remote sensing integration and machine learning:

Estimating Biomass Using Satellite Data: Combining Synthetic Aperture Radar (SAR) backscatter with LiDAR-derived canopy height models and multispectral vegetation indices allows for biomass to be estimated across an entire forest landscape with spatial continuity.

Machine Learning models trained on field inventory plots translate remote sensing signals into estimated carbon density with confidence limits. **Detecting Change on a Continuous Basis:** By using freely available satellite imagery with a 10-meter resolution (Sentinel-2, Landsat), it is possible to track forest disturbance on a near-daily basis throughout the world.

Agentic systems continuously process incoming imagery to automatically detect deforestation, selective logging, fire, and damage due to storms; generate real-time alerts for these disturbances; and provide the information necessary for regulators to respond quickly and invalidate carbon credits. **Predicting Growth Rate of Trees:** LSTM (Long Short Term Memory) neural networks trained on time series of satellite observations, climate data, and stand inventory data can predict growth rates of trees under multiple future climate scenarios, assisting with baseline projections and providing developers of carbon projects with better financial forecasting. **Automating Registry Tracking:** Agentic systems allow for direct integration with the APIs of carbon registries (Verra, Gold Standard, ACR), thereby automating requests for credit issuance, submission of monitoring reports, and submission of verification data packages. Automating these business processes will significantly reduce the administrative burden associated with managing carbon projects.

5.4 Economic Implications for Forest Owners and Buyers

Agentic artificial intelligence has economic implications for forest carbon flow. Agentic AI has a positive economic impact on the supply and demand side of the forest carbon market. Specifically, automatically brokering MRV for agricultural landowners and developers will decrease the annual monitoring cost (USD \$5–\$15 per hectare). People should expect to have their monitoring costs significantly reduced with automated monitoring (e.g., >USD \$1–\$3/hectare). Other costs associated with expecting to have all of your landowners be able to count on being participating viable in carbon markets is to have half a million dollar project (0.5 million hectares) and a half-million property (0.500 hectare)) will mean that half of all landowners will be able to participate in the market as a result of having an automated monitoring capability. An example of how large a purchase will be made in terms of the scale of the market is Microsoft's January 2025 announcement that they will purchase 3.5 million carbon credits to offset the emissions associated with their AI development. The availability of a machine verified monitoring data source along with continuous satellite monitoring records will provide a rapidly expanding body of information that will be used by buyers for their due diligence and will address any market integrity issues resulting from the reduction of the due diligence burden by buyers. The type and amount of monitoring data will allow buyers to determine credit quality/confidence to a much greater degree than ever before.

VI. IMPLEMENTATION FRAMEWORK

6.1 Technical Architecture Requirements

When deploying agentic AI in consulting and environmental sectors there are several architectural components to pay careful attention towards including:

- 1. Data Infrastructure** – Cloud data lakes (like AWS S3, Azure Data Lake and Google Cloud Storage) must be set up to store both *structured* & *unstructured* information in *Terabytes*. Data Ingestion pipelines should be built using *Kafka* (or *Pulsar*) for low-latency streaming parameterized from the Enterprise Systems to the Data Lake. All agent ingest (inputs) must meet reliability thresholds via *Data Quality Monitoring* tools (using *Great Expectations* or *Monte Carlo*), that use *automated validator* monitors.
- 2. Agent Orchestration Layer** – In order for specialized agents cooperating on complex problem domains to do so, *Multi-agent Frameworks* (like *Langchain*, *AutoGen* and *CrewAI*) provide the cooperation System Architecture that breaks the Task down into *Separate/Relationship Definitions*, assigns the agent(s) responsible for the task, aggregates the results of the inputs from the respective agent(s) who contributed to that Task input, and then determines what actions to take to resolve the situation when there are inconsistencies between the outputs from the Agent contributing to the tasks.
- 3. Processing Pipelines** – The core functional elements of the system include *Domain Specific NLP Pipelines* to perform Regulatory document parsing, *Anomaly Detection Models* (for Audit Data), *Geospatial Processing Pipelines* (for Satellite Data) & *Emission Calculation Engines*.
- 4. Delivery / Interface Layer** – This layer includes information provided to the User includes *Dashboard(s)*, *Alerts*, *Automated Reporting Pipelines* and *API Integration* to the Client's System, *Role Based Access Control* and *Explainability of Agent(s)* Output to allow an *Human Based Decision Maker(s)* ability to Investigate, Validate or Overrule Agent(s) outputs.

6.2 Change Management and Human-in-the-Loop Design

In the field of professional services, utilizing machines and humans together requires design considerations for their collaboration. The move from human-delivered consulting to

consulting augmented by machines (Agentic AI) does not simply replace existing tasks; it requires a redesign of roles, skill development, and culture change.

Redefinition of Roles: Historically, junior consultants have been the primary performers of tasks such as gathering and analyzing data, and creating an initial draft of reports. However, as Agentic AI has the ability to take over many of these tasks, leading firms have looked to redirect junior consultants to other responsibilities, including systems oversight, quality assurance, client communication, and more complex judgement functions requiring the judgment of an experienced consultant.

Trust Calibration: Professional consultants and clients need to establish calibrated trust in the output of Agentic AI. Therefore, they must utilize both the existing rules of the consulting profession with respect to accepting and or rejecting deliverables generated through Agentic AI systems. A necessary precursor to establishing calibrated trust is to have tools for explaining how an output was derived from the model and having documented evidence of how the model reached a conclusion.

Accountability and Liability Framework: For example, in regulated industries (auditing, financial advice, environmental compliance), developing clear accountability frameworks is critical. Firms must have clear policies and procedures that articulate who has liability for an error made by an Agentic AI system and how the professional responsible for the work the Agentic AI system performed will be held accountable for errors the Agentic AI system made while performing tasks that the professional historically owned.

6.3 Data Governance and Privacy

Robust data governance frameworks must be employed when using agentic systems to process sensitive financial, regulatory, and environment-related data. A few of the key requirements pertaining to the data governance framework include data lineage tracking (allowing an agent to trace any results produced back to the dataset that it originated from), implementing access controls to limit the accessibility of data based on an agent's assigned function, implementing privacy-centered computation techniques when working with sensitive supply chain data, and establishing model governance protocols addressing training dataset quality, model validation, as well as version management of the model.

VII. FUTURE TRAJECTORIES

7.1 Foundation Model Advancement

The agentic AI landscape is rapidly evolving as foundation models (large pretrained language models and multimodal models) are constantly evolving, which are the core cognitive models found within agentic systems. Increasing the number of parameters being trained, increasing the breadth of the training datasets, and improving the integration of multimodal data are expanding the capability of agents to exhibit richer and more complex behaviors.

References

1. MarketsandMarkets. (2025). Agentic AI Market by Component, Application, End-User and Region – Global Forecast to 2032. MarketsandMarkets, Pune, India.
2. Mordor Intelligence. (2025). Agentic AI Market – Growth, Trends and Forecasts (2025–2031). Mordor Intelligence, Hyderabad, India.
3. Precedence Research. (2025). Agentic Artificial Intelligence Market (By Component, Application, End-use, Region) – Global Industry Analysis and Forecast, 2024–2034. Precedence Research.
4. Fortune Business Insights. (2025). Agentic AI Market Size, Share and Industry Analysis, 2024–2034. Fortune Business Insights.
5. Dimension Market Research. (2025). Global Agentic AI Market – Analysis and Forecast to 2034. Dimension Market Research.
6. [Additional market study] (2024). Global Agentic AI Market Size, Share, Trends, and Forecast to 2034. (Source used for: “USD 5.43 billion in 2024” and “USD 236.03 billion at 45.82% CAGR.” Replace with the exact firm name and title you actually used.)
7. Global Market Insights (GM Insights). (2024). Voluntary Carbon Market Size, Share and Trends Analysis Report, 2024–2034. Global Market Insights.
8. Grand View Research. (2024). Voluntary Carbon Market Size, Share & Trends Analysis Report by Type, End-use, Region, and Segment Forecasts, 2024–2030. Grand View Research.
9. Roots Analysis. (2025). Voluntary Carbon Credit Market, 2025–2035. Roots Analysis.
10. Climate Focus. (2024). Voluntary Carbon Market 2024 Review. Climate Focus.
11. Association of Certified Fraud Examiners (ACFE). (2024). Report to the Nations: 2024 Global Study on Occupational Fraud and Abuse. ACFE.
12. European Parliament and the Council of the European Union. (2022). Directive (EU) 2022/2464 of the European Parliament and of the Council of 14 December 2022 as regards corporate sustainability reporting (CSRD). Official Journal of the European Union.
13. U.S. Securities and Exchange Commission (SEC). (2024). The Enhancement and Standardization of Climate-Related Disclosures for Investors. Final Rule. U.S. SEC, Washington, DC.
14. Task Force on Climate-related Financial Disclosures (TCFD). (2017). Recommendations of the Task Force on Climate-related Financial Disclosures. TCFD.
15. World Resources Institute (WRI) & World Business Council for Sustainable Development (WBCSD). (2004). The Greenhouse Gas Protocol: A Corporate Accounting and Reporting Standard (Revised Edition). WRI & WBCSD.
16. IPCC. (2006). 2006 IPCC Guidelines for National Greenhouse Gas Inventories, Volume 4: Agriculture, Forestry and Other Land Use. Intergovernmental Panel on Climate Change.
17. GOFC-GOLD. (2015). A Sourcebook of Methods and Procedures for Monitoring and Reporting Anthropogenic Greenhouse Gas Emissions and Removals from Deforestation, Forest Degradation and Other Forest Land Uses (GOFC-GOLD Report version 2.0). GOFC-GOLD Land Cover Project Office.
18. UNFCCC. (Various years). REDD+ Methodological Guidance and Decisions under the UNFCCC. United Nations Framework Convention on Climate Change.
19. Verra. (Various years). Verified Carbon Standard (VCS) Program Rules and AFOLU Methodologies. Verra.
20. Gold Standard Foundation. (Various years). Gold Standard for the Global Goals – Land Use & Forests Methodologies. Gold Standard.

21. Microsoft. (2025). Microsoft announces purchase of 3.5 million carbon credits to address emissions from AI development. Microsoft press release, January 2025. (Replace with exact title/date of the press release you used.)
22. Tata Consultancy Services (TCS). (2025). Agentic AI in Audit and Assurance: Transforming the Audit Value Chain. TCS Whitepaper.
23. ISACA. (2025). The Growing Challenge of Auditing Agentic AI. ISACA Industry News article.
24. The Institute of Internal Auditors (IIA). (2025). Agentic AI in Internal Auditing. All Things Internal Audit Podcast/Briefing.
25. Fieldguide. (2025). How to Use Agentic AI to Support Integrated ISO Audits. Fieldguide Resource Article.
26. KPMG International. (2026). AI and Technology in Audit. KPMG Audit Insights.
27. Salesforce. (2026). How Our AI Agents Save Us Time — And Scale Sustainability. Salesforce Blog.
28. 8allocate. (2026). Top 50 Agentic AI Implementations: Use Cases to Learn in 2026. 8allocate Blog.
29. Digiqt. (2025). AI Agents in Carbon Credits: Proven Wins & Pitfalls. Digiqt Blog.
30. DeBales AI. (2025). Automated ESG Reporting: AI Agents for Real-Time Carbon Optimization. DeBales Blog.
31. Arion Research. (2025). Agentic AI for Sustainability: Can Autonomous Agents Act as ESG Co-pilots? Arion Research Blog/Whitepaper.
32. Morgan, P. (2025). Auditing in the Age of Agentic AI – Why ITGC Will Never Be the Same. LinkedIn Article.
33. Breunig, M. M., Kriegel, H.-P., Ng, R. T., & Sander, J. (2000). LOF: Identifying density-based local outliers. Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data, 93–104.
34. Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. 2008 Eighth IEEE International Conference on Data Mining (ICDM), 413–422.
35. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
36. Apache Software Foundation. (2024). Apache Kafka Documentation. Version/current. Available at: <https://kafka.apache.org/>
37. Apache Software Foundation. (2024). Apache Pulsar Documentation. Version/current. Available at: <https://pulsar.apache.org/>
38. Superconductive. (2024). Great Expectations Documentation. Available at: <https://greatexpectations.io/>
39. Monte Carlo Data. (2024). Monte Carlo Data Observability Platform – Product Documentation. Available at: <https://www.montecarlodata.com/>
40. LangChain. (2024). LangChain Documentation and API Reference. Available at: <https://python.langchain.com/>
41. Microsoft. (2024). AutoGen: Enabling Next-Generation Large Language Model Applications via Multi-Agent Conversation. Microsoft/AutoGen Documentation & GitHub.
42. CrewAI. (2024). CrewAI Framework Documentation. Available at: <https://www.crewai.com/> (or official docs/GitHub you used).
43. Messier, W. F., Glover, S. M., & Prawitt, D. F. (latest ed.). *Auditing and Assurance Services: A Systematic Approach*. McGraw-Hill Education.
44. Ecosystem Marketplace / Forest Trends. (Various years). State of the Voluntary Carbon Markets series. Forest Trends Association.
45. Trove Research. (Various years). Global Voluntary Carbon Market Outlook and Forecasts. Trove Research.
46. World Bank. (2016). Forest Carbon Partnership Facility: Carbon Fund Methodological Framework. World Bank.
47. Arion Research / Other. (If you used any specific article for MRV cost ranges USD 5–15/ha/year, insert here with full bibliographic details.)
48. DeBales AI / Other. (If you used any additional MRV automation case study for agriculture/forests, specify here.)
49. World Resources Institute. (2020). GHG Protocol for Project Accounting. WRI.
50. Any additional primary source for “agentic systems valued at USD 5.43B in 2024, forecast 236.03B”: give exact firm, report title, year.