




Our Responsible, Ethical, and Trustworthy AI Principles

Why do we have these principles?

Even in a heavily regulated industry, insurance businesses must go above and beyond what's expected to earn customer trust.

At Bevaya, our singular focus is supporting insurance companies in delivering on their customer promises while keeping the trust and security of our AI at the forefront of everything we do.

For our customers, we've developed a comprehensive framework of eight core principles for responsible, ethical and trustworthy AI.



1 Commitment to Transparency and Accountability

Principle:

Ensure that our AI models and algorithms are transparent and understandable.

Explanation:

Since our founding in 2018, we've committed to thoroughly documenting the advanced models we've developed, trained, and fine-tuned, including Machine Learning, Deep Learning, and Large Language Models. This transparency ensures our customers clearly understand how data is identified, classified, and extracted to provide insight into how their fine-tuned models will continue to improve and learn over time.

Approach:

Bevaya prioritizes transparency by ensuring users are informed about their interactions with our models, including identifying the specific models in use.



We provide **clear visibility into our data sources, algorithms, training methodology, model safeguards**, and other critical aspects of our AI systems.

Our internal governance policies, including regular reviews and audits, ensure our AI adheres to strict ethical standards.

2 Explainability Is a Top Priority

Principle:

Build AI that delivers clear and understandable insights into decisions and actions.

Explanation:

Implement practices and tools that ensure AI decisions can be understood and validated by all parties.

Transparent and interpretable AI processes let users, stakeholders, and regulators understand how decisions are made to drive accountability, fairness, and effective human-AI collaboration.

Approach:

We build AI models from the outset, using interpretable algorithms where possible—and we've created tools to translate complex model outputs into user-friendly explanations. An example is our Confidence Score feature, which delivers output confidence reporting in the workflow, informing users when document processing can continue straight-through, or when the model needs human review.



We collect user feedback about the **clarity and usefulness of explanations** generated by our AI through our patented Human-in-the-Loop [HITL] capabilities. This feedback is used to continuously improve our products' ability to provide understandable and usable insights to customers.

3 Safeguarding Privacy and Data Security

Principle:

Respect and protect the privacy of individuals and companies by ensuring maximum data security while minimizing data usage.

Explanation:

Our philosophy is simple:

Your data belongs to you, not us, and we will keep it that way.

Integrity and security of personal data are paramount.

We respect your data as you would because it's your data.

We live up to this belief by applying the most stringent security measures to protect sensitive information.

Additionally, we adhere to the principle of data minimization, only collecting and using data necessary for the AI's functionality and continuously evaluating data retention.

Approach:

Our services run on Microsoft Azure and Amazon AWS cloud platforms. Unlike public AI solutions (e.g., OpenAI, Meta, LLaMA, and others), your data is never exposed to the open internet and is fully secured in our environment.

We do not employ any on-premise routers, load balancers, DNS servers, or physical servers. We use advanced end-to-end encryption for data in transit and the most robust encryption and security practices for data at rest. At every stage of AI development we integrate privacy considerations, from data collection to model deployment, ensuring that privacy protection is a fundamental part of the system architecture.

We do not build or train models in a way that allows them to learn, memorize, or reproduce customer data—and we do not use any of our customers' data to train models built for use by another customer.



We conduct **continuous security evaluations and updates to protect against data breaches** and ensure that all data handling complies with relevant regulations, such as GDPR or CCPA.

4 Embedding AI Reliability and Safety

Principle:

Develop robust, reliable, and safe models for all users and the customers they serve.

Explanation:

Bevaya's AI systems are thoroughly tested and continuously monitored to ensure reliability of output. We implement mechanisms to eliminate potential risks and unintended consequences, to create insurance AI solutions users can depend on when the stakes are highest.

Approach:

Bevaya is a business founded by experienced insurance operators; as such, our policy is to ensure our tech supports humans making underwriting or claims decisions. Our models are trained to predict based only on document contents.

We perform extensive testing of AI models under different scenarios to ensure consistent performance, reliability, and accuracy across a full range of conditions. Once in production, we monitor models in real-time to detect and address anomalies, errors, or drifts in model performance as they occur.



Bevaya's AI is built with fail-safes/fallback procedures that activate when it encounters situations it cannot handle reliably. Our patented Human-in-the-Loop (HITL) **functionality ensures accuracy through human intervention when needed**, allowing customers to control throughout and review AI decisions in real-time.

5 Removing Bias and Prioritizing Fairness and Inclusivity

Principle:

Design AI systems that are fair, equitable, and inclusive, avoiding biases that could lead to discrimination.

Explanation:

AI should not reinforce existing biases or create new forms of discrimination. We are committed to regularly auditing our AI models to detect and mitigate bias, ensuring that our solutions are inclusive and equitable for all users, regardless of their background, demographic, or other personal characteristics.

Approach:

Bevaya doesn't make judgments, nor do we use input data (e.g., chatbots) that potentially introduce bias. We regularly audit, test, and fine-tune our models to mitigate potential biases, especially those related to race, gender, age, or socioeconomic status. We seek to engage diverse teams and stakeholders in the AI development process to identify and eliminate potential biases early on.



We actively work to reduce biases and **regularly audit, test, and fine-tune** our models.

6 Model Testing, Reproducibility, Robustness, and validation

Principle:

Establish rigorous processes for testing, reproducing and validating AI models to ensure they are robust, reliable and perform consistently in real-world scenarios.

Explanation:

Bevaya's AI is built, trained, tested, and fine-tuned by our in-house team with 50+ years of combined experience. Our adherence to this principle is critical to minimizing risks and ensuring our AI systems deliver excellent results in all conditions.

Approach:

Our testing protocols evaluate AI models under diverse scenarios, including edge cases and stress tests, to ensure reliable performance in real-world conditions. We use standardized tools and methodologies to ensure that AI models can be reliably reproduced in different environments, providing consistency in results across various deployments.



We continuously validate models using new data and real-world scenarios, **regularly updating them to reflect changes in data patterns or underlying assumptions** to maintain accuracy and robustness.

7 Establishing Data Governance and Compliance

Principle:

Implement robust data governance frameworks and ensure compliance with all relevant legal and ethical standards.

Explanation:

Effective data governance is crucial for maintaining the integrity, security, and privacy of data used in AI systems. Compliance with legal and ethical standards, such as GDPR or CCPA, ensures that data practices are responsible and protect individuals' rights. This principle underpins the trustworthiness of AI systems by ensuring that data is handled in a secure, ethical, and compliant manner.

Approach:

Bevaya performs a yearly third-party review for SOC 2 Type 2 attestation (report available on request). We are compliant with NIST, HIPAA, CCPA, GDPR, and 23 NYCRR Part 500 standards to protect all Personal Financial Information (PFI), Personal Health Information (PHI), and Personally Identifiable Information (PII).



Our comprehensive data governance policies outline responsibilities, standards, and practices for data management across the organization. **We conduct frequent audits to ensure data practices comply with legal regulations and industry standards, identifying and addressing any gaps or risks in data handling processes.** Bevaya's applications and infrastructure are regularly assessed for security vulnerabilities by third-party security services.

8 Fostering Human-AI Collaboration

Principle:

Promote collaboration between human experts and AI, leveraging the strengths of both.

Explanation:

AI works best by enhancing human expertise. Our systems are built to support and augment human decision-making, creating AI that empowers users. Making AI that collaborates is the crux of responsible AI that puts human judgment in command of complex, high-stakes decision-making.

Approach:

We configure our AI to support and enhance human decision-making rather than replace it.

Bevaya's patented Human-in-the-Loop (HITL) capabilities give underwriting and claims experts control over low-confidence extraction situations. Users can set their confidence thresholds and create default alerts that prompt human intervention with output.

As part of customer onboarding, we provide extensive training and ongoing support to help users understand how to collaborate most effectively with their AI-powered tools. This includes understanding systems' limitations and best practices for use in their decision-making processes.



Our AI has established mechanisms to collect users' feedback on AI performance and use this feedback to continuously improve the system. **This ensures that our solutions evolve in a way that remains aligned with our customers' unique needs.**

FAQs

Which compliance frameworks does Roots adhere to?

ISO 27001, HIPAA, SOC 2 Type 2, OWASP, CCPA, GDPR, and 23 NYCRR Part 500.

How is customer data protected during transmission and storage?

All data is encrypted with 256 AES at a minimum during transmission and at rest.

Where is customer data stored?

In cloud provider data centers in the US and Europe (depending on customer requirements and appropriately segregated by customer).

How long is the data stored?

Dependent on the customer's needs and contract requirements.

What measures are in place to prevent unauthorized access to customer data?

Bevaya uses Role-Based Access Control (RBAC), per service/users, along with standard network and firewall controls in place.

How often is data backed up, and what is the backup retention policy?

Services requiring backup in the production environment are performed daily/weekly, with retentions based on contractual requirements. All data is stored across multiple data centers.

Does Bevaya perform regular security audits and testing?

Bevaya conducts third-party quarterly assessments on public services, third-party annual penetration testing on applications, and quarterly security reviews across our control environments.

What happens in the event of a security breach?

Bevaya rapidly initiates and deploys a prioritized set of remediation incident responses, including reviewing the scope of the breach, performing Root Cause Analysis (RCA) and engaging law enforcement as needed. Communication around RCA will be delivered promptly to the affected customer at the time of the breach and thereafter as needed.

Can customers control who has access to specific data and features within their organization?

Your system administrators can control customer team members' access types.

Is customer data shared with other customers or vendors?

No.

Additional Links

[Privacy Policy](#)

[Terms and Conditions](#)

[Security Standards and Practices](#)

[Contact Us](#)



About Bevaya

Bevaya is the AI Agent platform built for insurance, by Roots Automation, Inc. Our AI Agents read, analyze, and recommend across underwriting, claims, and policy servicing. They're powered by InsurGPT™ and proven across 115+ production deployments.

© 2026 Roots Automation, Inc.

bevaya.ai