Advisor.AI - Information Security

At Advisor.AI, data privacy, integrity, and security are of the utmost importance. This document provides a detailed overview of the measures we have in place to ensure the security of the Advisor.AI platform and to safeguard our users' information.

As a pioneering AI-powered advising platform, Advisor.AI is committed to supporting students, educators, and institutions with personalized and data-driven tools to enhance student success. With this commitment comes a responsibility to protect sensitive user data. We have implemented various security measures, including advanced encryption protocols and strict access controls, to ensure the safety of all user information.

Advisor.AI conducts regular security audits and assessments to identify potential risks or vulnerabilities, allowing us to take proactive steps to secure our system and services. With these robust measures in place, Advisor.AI provides its customers and users with peace of mind, ensuring that sensitive data is handled with care and remains fully protected.

Our security and ethical AI team comprises senior engineering experts, ethical AI practitioners from Fortune 500 tech companies, legal professionals specializing in data privacy, and leaders from various public and private universities, all working together to ensure that robust security and ethical AI best practices are followed and implemented.

All of our industry-leading security and responsible AI services are included as part of the Advisor.AI platform at no additional cost. This sets us apart as the only comprehensive AI advising solution with no hidden fees related to security, AI model training, and implementation.

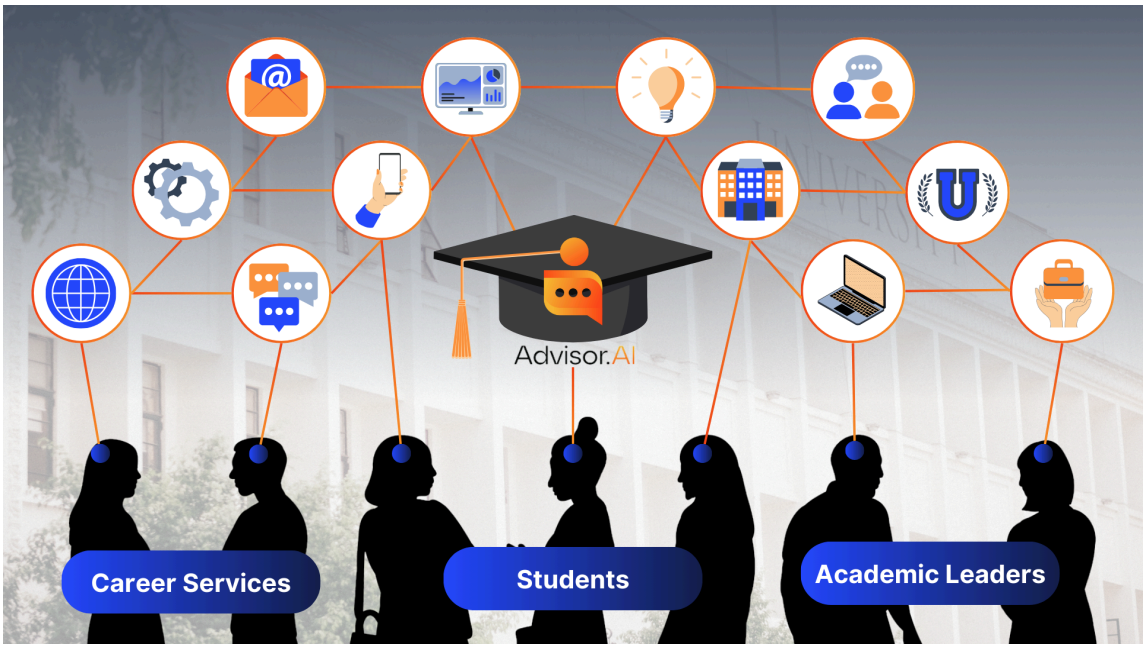

Platform Security

Access Control Policy

Access to data within the Advisor.AI platform is strictly managed through robust access controls, ensuring that all users—whether students, advisors, or administrators—have the appropriate permissions based on their roles. Advisor.AI implements a closed-system approach to data entry and access, limiting exposure to and entry of sensitive information into the system. Our predefined roles and permissions adhere to the least-privilege principle, making it easy for administrators to assign the correct access levels for each user while safeguarding data privacy and system integrity.

Secure Authentication

Access to administration interfaces is encrypted using industry-leading protocols like HTTPS and TLS 1.3, with AI models and data privacy safeguarded through end-to-end security measures. Advisor.AI administrators handle university registration and de-registration.

Data in Transit

Advisor.AI uses industry-leading best practice encryption schemes (HTTPS and TLS 1.3) to encrypt data in transit between the platform and users (students, advisors, and administrators). We disable deprecated protocols and weak cipher suites.

Data at Rest

Customer data at rest is encrypted by default using AWS service-side encryption with AWS-managed keys. All storage layers use AES-256 encryption.

Availability

Critical system components are backed up across multiple, isolated locations known as Availability Zones. Each Availability Zone is engineered to operate independently with high reliability. Availability Zones are connected to enable us to easily architect applications that automatically fail over between Availability Zones without interruption.

Operations Security

Change Management

All production changes are tracked, reviewed, and approved by someone other than the author. Implementation is achieved through controlled deployment processes, which are monitored after release and documented with a backout plan. Emergency changes follow expedited procedures with mandatory post-implementation review.

Development Process

Our development process follows a secure SDLC aligned to OWASP and SOC 2. All code and infrastructure-as-code changes are stored in version control; merges require peer review by a non-author and must pass CI gates. High-risk changes undergo threat modeling and, when appropriate, DAST and security sign-off. Secrets are stored in AWS Secrets Manager/Parameter Store and never committed to repositories. Environments are separated (dev/test/staging/prod) with least-privilege access and MFA, and critical vulnerabilities are remediated within defined SLAs. The process integrates with change management for approvals, rollout plans, and rollbacks.

Compliance

Our compliance program is built on the AWS Shared Responsibility Model and aligned with recognized standards (SOC 2 Trust Services Criteria and the NIST Cybersecurity Framework). We maintain documented policies and risk assessments; enforce least-privilege access, encryption, change management, secure SDLC, and vulnerability management; conduct employee security and privacy training; and perform independent security testing. We continuously monitor control effectiveness via automated checks and evidence (tickets/PRs, scan reports, CloudTrail/Config records) to support audits and customer reviews.

Responsible & Ethical AI

Responsible AI is the practice of developing, using, and deploying artificial intelligence (AI) systems in a way that is ethical, transparent, and accountable. The goal of responsible AI is to ensure that AI technologies are aligned with human values and promote the well-being of society and individuals. Our team of legal experts, AI practitioners, and university leaders meet regularly to assess our processes in relation to industry-leading frameworks such as SAFE.

SAFETY

At Advisor.AI, we prioritize data security through robust encryption practices. All training data and model outputs are encrypted both in transit and at rest, ensuring the highest levels of protection throughout the entire data lifecycle. We employ industry-standard encryption protocols to prevent unauthorized access and safeguard sensitive information. Additionally, role-based access controls are implemented, allowing only authorized personnel to interact with model training and inference APIs. These security measures ensure that the platform remains resilient against potential threats, maintaining the confidentiality of our AI-driven insights.

Some of the key features for safety include:

- Redacting sensitive information (PII) to protect privacy
- Blocking inappropriate content with custom word filters

- Detecting hallucinations in model responses using contextual grounding checks
- Bringing a consistent level of AI safety across our application via continuous monitoring
- Filtering harmful content based on our responsible AI policies

ACCOUNTABILITY

We foster accountability through a system that ensures active collaboration between students and advisors. Our platform is designed to prevent over-reliance on AI by empowering human oversight. Advisors can engage with students directly through the platform, while students have access to advisors and mentors via the app, ensuring that human feedback remains integral to academic and career decisions. This collaborative approach ensures both our AI systems and users are accountable for driving meaningful outcomes.

FAIRNESS AND TRANSPARENCY

Fairness and transparency are foundational to our AI development process. We proactively mitigate bias by rigorously testing our proprietary models before deployment, ensuring that they deliver equitable results across diverse student profiles. Continuous monitoring and weekly model audits help us maintain fairness and ethical standards. Transparency is also key: we make our AI-driven recommendations clear and understandable to users, supplemented by comprehensive FAQs that explain how decisions are made. This openness fosters trust and confidence in the recommendations powered by the platform.

EFFICACY

Advisor.AI is built to achieve measurable success outcomes, guiding students from pre-enrollment to career readiness. Our AI-driven insights work in conjunction with human support systems to enhance, rather than replace, personal advisor-student interactions. The platform's reliability ensures that students receive timely, accurate, and personalized support throughout their academic journey, helping them stay on track and succeed in their educational goals.

Technical Operational Measures

Environment Separation

Non-production environments are logically divided from the production environment.. Production data is never used in lower environments.

Backup

Our backup policy ensures that platform data on Advisor AI is replicated across multiple geographical locations. Our production databases are backed up and versioned every day. Those backups are kept for seven days. Backups are fully encrypted.

Logging and Monitoring

We centralize and continuously monitor application, infrastructure, and security events across our AWS environments. AWS CloudTrail captures account activity; CloudWatch Logs/Metrics/Alarms, VPC Flow Logs, ALB/ELB access logs, RDS engine logs, and S3 access logs feed a centralized log store for correlation and detection. Logs are time-synchronized, encrypted at rest using AWS-managed keys, access-controlled with least privilege, and retained per defined policy; integrity controls and change auditing help ensure tamper resistance.

Logical Access

Access to the Advisor AI's production infrastructure is restricted to specific members of the Advisor AI's engineering team, following the least-privilege principle.

Physical Perimeters and Location

Our platform is hosted in Amazon Web Services (AWS) facilities in the US. AWS data centers are housed in nondescript facilities. Physical access is strictly controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access the data center floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff. The AWS physical access points to server rooms are recorded by Closed-Circuit Television Cameras (CCTV). Images are retained in accordance with legal and compliance requirements.

Data Retention

Customer data is retained only as long as necessary to deliver the service, meet contractual obligations, or fulfill legal/regulatory requirements. By default, following contract termination, customer data is retained for up to three months. Retention and deletion are enforced using AWS-native controls, and backups and logs follow predefined retention periods. When the retention period expires—or upon verified customer request—data is securely and permanently deleted via cryptographic erasure and controlled purge processes, rendering it unrecoverable; decommissioned media sanitization is performed by AWS per NIST SP 800-88. Legal holds suspend deletion, and exports are supported prior to deletion when applicable. All retained copies remain encrypted at rest and protected by least-privilege access.