Data Compliance & Security

We go above and beyond legal requirements to protect all of your data. This document outlines additional precautions we take to keep your customers' data safe.

ISO 27001 compliant data centers

Our technical infrastructure is provided by Amazon Web Services, which maintains a number of globally recognized compliance certifications, to include Cloud Security Alliance, ISO 27001, PCI Level 1, SOC 2 and SOC 1. All of our services are cloud-based. We do not run our own infrastructure.

Our data services are hosted in Amazon Web Services (AWS) facilities across the EU. We do not have control over where our third-party services are hosted, for example, Google Analytics. All of our clients' data is hosted and processed using AWS.

Our infrastructure is spread across 3 AWS data centres (also known as availability zones). This adds redundancy to our system, as should one of the data centres fail unexpectedly, our services will continue to work.

Network Design

Our network design is based on a 3-tier structure, where the internal network is separated from the DMZ, and the DMZ is separated from the external network. Detailed information is available upon further request.

Firewalls

We use Amazon EC2 Security Groups which act as a virtual firewall that controls the traffic for our server instances. In addition, we use Amazon GuardDuty to perform threat detection. GuardDuty continuously monitors for malicious or unauthorized behaviour on our AWS accounts and workloads. It monitors for activity such as unusual API calls or potentially unauthorized deployments that indicate a possible account compromise. Additionally, it is used to detect potentially compromised instances or reconnaissance by attackers.

Firewall rules are reviewed on a daily basis, to detect rules that could create vulnerabilities and remove unused rules. We are also automatically notified when a new firewall rule is added and this is reviewed. Only senior engineers may update firewall rules.

Data storage and access

All of our client data is stored in the European Union.

Client data is stored in multi-tenant data stores. This means that we do not have individual data stores for each client. Should you wish to have your own dedicated datastore, please contact us and we can discuss your requirements.

In order to prevent one client from accessing another client's data, we have a number of low-level code checks that fail upon not being provided with the logged-in customer identifier. We employ automated testing prior to every code change being deployed on our production services. Additionally, we periodically perform code audits to prevent this from happening.

Internally, all database entity types have a client identifier field. All queries are required to provide a client identifier. This check has been implemented at a low-level. It ensures that one client cannot access the data of another client.

We run in a virtual private cloud (VPC). Our database services live inside our private subnet. This means that only servers in the public subnet can communicate with our database servers. All ports (besides HTTP (80) and HTTPS (443)) on servers living inside our public subnet have been restricted to whitelisted IPs defined in a security group. The whitelisted IPs are the addresses that we use internally and are inaccessible from sources outside our network. We only permit server access to public keys whitelisted on our servers. This prevents SSH server access from computer devices outside of our organisation.

Encryption

Your data is encrypted at rest in AWS S3 buckets, AWS RDS instances and our internally-managed databases. AES256 encryption is used by default via AWS' encryption services. Key management is handled by AWS KMS. This ensures the content is preserved and safe from prying eyes and manipulation.

All data sent to and from Traitly is encrypted in transit using state-of-the-art 256-bit encryption.

Our platform and API are SSL-only.

All communication between you and Traitly, that includes your data, traverses the Internet via encrypted HTTPS traffic using TLS v1.2. This encryption during communication ensures information cannot be read or manipulated by unauthorized third parties.

Code reviews and secure development lifecycle

We implement the “secure by design” philosophy, whereby security features are embedded in the product and our overall architectural design. This ensures that new and existing functionalities are free of vulnerabilities. In addition, we conduct code reviews through the implementation of a Secure Software Development Life Cycle (SSDLC) framework, where the code is reviewed by peers prior to being merged into our testing, staging and production environments. In addition, perform weekly team-level code reviews.

Personnel

Our engineering team includes people who have played significant roles in both startups and large organisations. They have experience building Internet-facing applications that house highly confidential and mission-critical data.

Incident response

In the event of a data security incident, all key personnel are requested to respond immediately. Those in charge of affected parts of our application and infrastructure are notified and assembled to address the incident quickly. Upon notification, incident resolution time is typically on the order of minutes.

Following a data security incident, a post-mortem analysis is performed. The outcome of our analysis is discussed internally and shared among the relevant personnel. The analysis includes actionable items to help make it easier to detect and prevent the occurrence of similar incidents in future.

Backup and redundancy

All mission critical systems are redundant. Our infrastructure components are deployed in at least three availability zones on AWS, minimizing disruptions caused by any failure and keeping your data available. Elastic Load Balancers are used to automatically split the load and segregate traffic from the Internet to all nodes of our frontend layer. Amazon Aurora is our backup manager. A complete client data backup is taken once every 6 hours. The backed up data is then stored on Amazon S3 in encrypted, where it is available for up to 7 days. You are free to download all your data to back it up off-site by fetching data via our API.

Access to client data

Access to client data is very restricted. We hand-pick and train engineers and support staff who, after your explicit permission, are permitted to help fix problems by accessing data that you authorize. These actions are recorded, audited and monitored using internal monitoring. Support staff may also require access to your data during the setup and onboarding process. This access is only granted following your explicit permission.

Data retention

All client data is retained only for as long as is necessary. At the end of the engagement, a client may choose to export data via our API. All sensitive data is then destroyed.

Distributed denial of service (DDOS) protection

Our APIs and web application are protected in multiple ways against denial of service attacks. AWS provides volumetric denial of service protection through AWS Shield and Elastic Load Balancing to ensure high availability. Amazon comes with a built-in network and security monitoring systems designed to provide increased protection against threats like Distributed Denial of Service (DDoS), Man in the Middle (MITM) attacks, password brute-force detection, and packet sniffing.

Build Automation

We deploy code on our production server tens, and in some cases, hundreds, of times per day. This enables us to respond to data security incidents with code changes within minutes.
We have a semi-automated deployment system, which requires us to peer-review all code changes before being deployed to our production servers. Code changes are reflected across all of our production servers within minutes. We use GitHub and AWS to help automate this process.

Authentication

Traitly is served 100% over HTTPS.

We use two-factor authentication (2FA) and stringent password policies across our own and third-party services we use. These include GitHub, AWS, Google, and Traitly.

Monitoring and Reporting

We use CloudTrail to log and monitor all events performed on our infrastructure. CloudTrail provides details about changes that are made, when they are made, and by whom. This data is periodically exported to S3. This data is also retained for 90 days.

Application

We use our own internal audit trail and logging system to log the actions of our own staff made on the Traitly platform. This includes data about when a staff member accessed a client account, the information they viewed, and any changes that were made. This is accompanied by a staff member identifier and the device used to access the data.

Vulnerability Management

We have a complete list of all security controls in place, e.g., security groups, firewalls, and IDses (AWS GuardDuty), and CloudFormation Templates (used for secure environment configuration). This list is used should we need to respond to a vulnerability alert.

We use a third-party to receive vulnerability alerts. We compare reported vulnerabilities with those listed in our inventory and control list. In the event that a vulnerability is identified, it is assigned a score, using the CVSS scoring system, and an owner. We have an internal SLA that stipulates deadlines for fixing vulnerabilities, while progress is tracked by tools and, if necessary, a post-mortem is arranged as a learning exercise for our engineers to improve code security. We share this responsibility among our in-house engineering team.

We apply the relevant patches. In addition, we take note of patches applied and update our inventory of systems and controls, where necessary. We seek to deploy patches without disrupting uptime or production.

Quality Assurance (QA)

We make use of Git, a version-control system for tracking changes to our code base. We use GitHub to review and share code commits among the engineering team. This is used as part of our code security review. All code changes are reviewed by at least two engineers. Code is not shipped to production until it has been successfully run and passes unit and integration tests on our development and staging environments, after which is must be reviewed. Following a successful review, code is merged with our main production branch.

Multi-factor authentication and SSO

Two-factor authentication and SSO, based on the SAML 2.0 standard, is available for clients. We encourage the use of two-factor authentication.

Data Processing Agreement

1. Purpose of processing

This Document ("Agreement") sets out the legal agreement between you, your directors, employees, contractors, agents and assigns, collectively the "Customer" and the "Company": AIBL TECHNOLOGIES LIMITED, an Irish incorporated entity with its registered offices at Unit 8, Crawford Commercial Park, Bishop Street, Cork Ireland ("Traitly") for compliance with General Data Protection Regulation (GDPR).

The Company is engaged by Customer to provide a dynamic knowledge base and support services logic and machine learning techniques (hereinafter - the Purpose), which requires the Company's access to the data about the employees of the Customer and data pertaining to the Customer.

2. Data processed

Data contained in the third-party platforms the Customer uses, as well as data generated within the Company’s platform (hereinafter together the Data).

3. Duration of processing

The Company shall process Data for no longer than required for the purpose of the Agreement between the Company and the Customer.

4. General Recipient's obligations

Company is obliged to:

- a) process the Data only to the extent such processing is needed for the purpose;
- b) ensure that its employees, directors and other officers having access to the Data within all the period of this Agreement (also after termination of their employment, contractual and other relations with the Company) are bound, whether via contract or statutory obligation to keep the Data confidential in accordance with the terms of this Agreement applicable to the Company;
- c) give access to the Data to a limited number of employees, directors and other officers, who need to know the Data for the Purpose;
- d) fully comply with Regulation (EU) 2016/679 (General Data Protection Regulation) and other applicable laws and regulations;

- e) process the Data only on documented instructions from the Customer;
- f) take all appropriate technical and organisational measures required under Article 32 of General Data Protection Regulation to ensure the security of Data;
- g) assist the Customer by appropriate technical and organisational measures, insofar as this is possible, for the fulfilment of the Customer's obligation to respond to requests for exercising the data subject's rights;
- h) notify the Customer immediately if it becomes aware of any breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorized disclosure of, or access to, Data and provide such further information as the Customer may reasonably require;
- i) assist the Customer in ensuring compliance with the obligations relating to the security of processing, data breach notification and data protection impact assessment, taking into account the nature of processing and the information available to the Company;
- j) make available to the Customer all information necessary to demonstrate compliance with the obligations laid down in this Data Processing Agreement and allow for and contribute to audits, including inspections, conducted by the Customer or another auditor mandated by the Customer.

The Company shall process Data for no longer than required for the purpose of the Agreement between the Company and the Customer.

5. Sub-processing

The Company can engage third parties to process the Data provided (i) this is required for the Purpose and (ii) it has obtained a specific written authorization of the Customer. Where the Company engages another processor for carrying out specific processing activities, it (i) bears full responsibility for the actions of third persons, to which it disclosed the Data, with regard to such Data and (ii) warrants and represents that the same data protection obligations as set out in this Data Processing Agreement shall be imposed on that other processor by way of a contract. In particular, the Company is obliged to ensure that such third persons:

- a) are informed on the confidential character of the Data provided to them;
- b) at the time of disclosure are bound via contract to keep the Data confidential substantially in accordance with the terms of this Agreement applicable to the Company;
- c) give access to the Data to a limited number of employees who need to know the Data for the purpose they were received for;

- d) do not disclose the Data to any other third person;
- e) use the Data purely for purposes for which it was provided;
- f) comply with confidential undertakings established by this Agreement for the Company, as they were the Company;
- g) return or destroy the Data once they are no more needed for the purpose.

6. Order of Precedence.

This Data Processing Agreement is an integral part of the License Agreement entered into upon registration with the Company's software by the Customer (hereinafter – the Agreement). Provisions of this Agreement prevail over other provisions of the Agreement in case of their contradiction.