Unbabel

THE COMPLETE UNBABEL SECURITY GUIDE

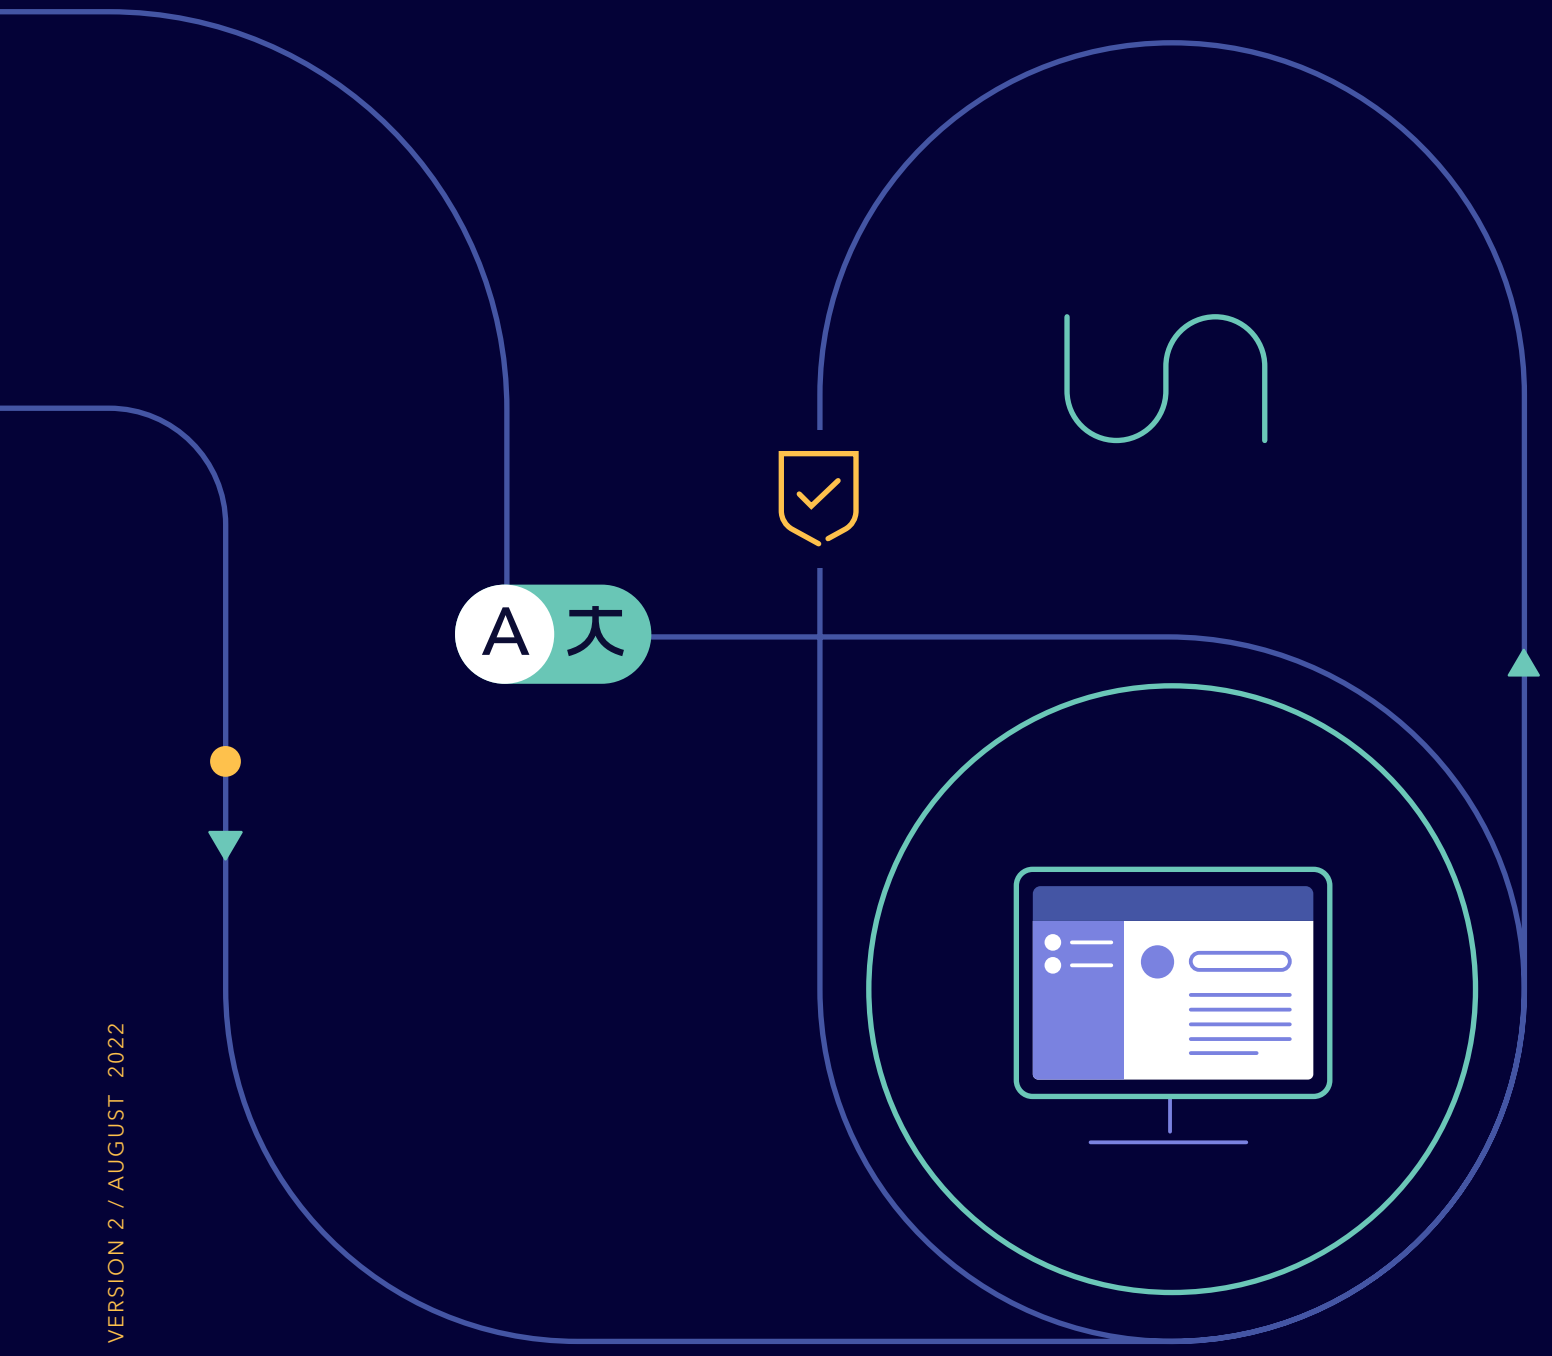

VERSION 2 / AUGUST 2022

Security at Unbabel

As a company focused on creating universal understanding, we at Unbabel believe that trust and transparency are key to building healthy relationships with our customers, partners, and other stakeholders. In particular, we recognize our responsibility in maintaining the confidentiality and privacy of communication between our customers and their customers as we translate. We believe strongly in the value of security and privacy and have incorporated best practices throughout our service. To that end, we have built Unbabel to be secure by design.

We know that our customers, who submit their own customers' communications into the Unbabel Platform for translation, must feel confident that neither their brand's reputation nor their own customers' sensitive data is at unnecessary risk. Since we deploy AI translation engines as well as a community of editors who review translations, we work hard to ensure privacy and security on the technological front and the human one as well, which we'll detail in this guide.

Rest assured that your data is protected with Unbabel. We leverage robust data encryption and anonymization mechanisms to ensure that all data is securely processed through our platform. Our solution is also ISO 27001 certified, as well as GDPR and CCPA compliant, as Unbabel is strongly committed to industry best practices in information security with the protection of customer data in mind.

Our team is also happy to answer questions that may arise, so please feel free to get in touch via security@unbabel.com if you require further clarification or insight.

Thank you for reading,
The Unbabel Security Team

Table of Contents

About Unbabel

Unbabel eliminates language barriers so that businesses can thrive across cultures and geographies.

The company's language operations platform blends advanced artificial intelligence with human editors, for fast, efficient, high-quality translations that get smarter over time. Unbabel integrates seamlessly in any channel, so agents can deliver consistent multilingual support from within their existing workflows. Making it easy for enterprises to grow into new markets and build customer trust in every corner of the world.

Based in San Francisco, Calif., Unbabel works with leading customer support teams at brands such as Facebook, Microsoft, Booking.com, and Under Armour, to communicate effortlessly with customers around the world, no matter what language they speak.

For more information visit **WWW.UNBABEL.COM**

For additional relevant information on security and privacy, please see our **PRIVACY POLICY**

Organizational Security

SECURITY TEAM

The Lead Security Engineer heads up information security at Unbabel, in close collaboration with our VP of Engineering. A Security and Privacy Committee, headed by the VP of People, is responsible for ensuring the implementation of the Information Security Policy, which is supported by a working group of around 15 people, part of which site reliability engineering (SRE), legal and HR, which all collaborate together to uphold security at Unbabel.

Unbabel also has a Data Protection Officer (DPO), who (i) monitors compliance of data processing with applicable standards, (ii) is a point of contact with the Data Subjects to clarify questions regarding the processing of your data by Unbabel, (iii) cooperates with the data protection authority, and (iv) provides advice about Unbabel's obligations regarding privacy and data protection. If after reviewing this document you have any inquiries or concerns about the handling of data or about our privacy practices generally, please contact our DPO at data-protection-officer@unbabel.com.

The Tech Ops team is responsible for the security processes and protection of Unbabel's information with regards to end-user computing and the provisioning and de-provisioning of users on corporate services. The Office Management team is further responsible for securing Unbabel buildings and other assets.

The site reliability engineering (SRE) team is responsible for the security and monitoring of Unbabel production and test environments and access to those environments through secure VPNs. We also have an incident management process implemented at Unbabel, coordinated with the overall security team and delegated to the relevant teams for incident response and post-mortem analysis.

Apart from these people, everyone at Unbabel has an information security responsibility and is trained regularly on policies and best practices.

SECURITY GOVERNANCE

We have a formalized Information Security Policy and are continuously working towards the goal of certifying an information security management system (ISMS) based upon the ISO 27001 standard, the goal of which is to ensure the confidentiality, integrity, and availability of Unbabel's information systems and your data. Our security policies are reviewed annually at a minimum.

SECURITY AWARENESS TRAINING

We have mandatory security awareness and training programs in place for all employees and temporary staff. This includes yearly training sessions for all staff on security hygiene best practices (covering a wide range of topics such as awareness of the surrounding environment, phishing, social engineering, remote working) and more focused domain-based training for specific teams: secure development and engineering practices, auditing and vulnerability management, hardening best practices, etc.

All of our staff understand their responsibilities to keep customer data secure and undergo regular training on Unbabel policies and general best practices.

HUMAN RESOURCES

All employees and temporary staff are required to undergo background checks where permitted by local laws. All employees and staff are required to sign confidentiality agreements. Employees are required to work to ensure the confidentiality of our customer data and ensure that Unbabel meets its legal, regulatory, and contractual responsibilities.

PHYSICAL SECURITY

Unbabel is focused on maintaining the security of its facilities, staff, equipment and data. Access control with the use of electronic keycards and video surveillance is implemented in all of our offices.

SECURITY IN REMOTE WORKING

Unbabel strives to view security holistically, identifying any risks to the organization and/or our customers. Our teams working in a telework setting are made aware of specific threats that may rise from these external environments through our security awareness program.

We secure access and communications to our infrastructure with the use of Unbabel's VPN. Moreover, our employees' devices are controlled and monitored continuously through our mobile device management (MDM) tool.

INCIDENT RESPONSE PLAN

Unbabel has a formalized incident response plan (IRP) that applies to all security incidents, technological or physical. We define an information security incident as a "single or a series of unwanted or unexpected information security events that have a significant probability of compromising business operations and threatening information security," per ISO/iEC 27000:2009. The IRP prescribes the functions and responsibilities for each relevant team member during an incident. This includes engineering, legal, and security team members.

Should a security incident occur that affects customers, Unbabel will notify the affected customers per our legal obligations, including those detailed in the EU GDPR and CCPA regulations.

Product Security

Security is always a key requirement from the beginning of any product design process at Unbabel. We maintain high standards in our testing and follow a secure development life cycle. We are dedicated to the continuous improvement philosophy. We also conduct regular third-party testing, as well as maintaining close engagement with the ethical hacker community.

ACCESS CONTROL

Duties are segregated and data access is granted on a “need to know” basis only. Access to data is provided only for those supporting individual customers in delivering their service. All users are authorized to access only information relevant to their activities; no other information can be accessed, and access is revoked once the need to access said information ceases or the employee relationship with Unbabel ends by resignation or termination.

From time to time, access may be required for additional resources. This is provided on a temporary basis, based on a “least-privilege” principle and revoked when access is no longer required. All additional access requests must be approved by the VP of Engineering.

SECURE TRANSLATIONS WORKFLOW

Since translations are at the heart of the Unbabel service, we believe in providing our customers and partners with a clear understanding of how we maintain a secure workflow. Communications are passed securely through our AI translation engines and then, in some cases, to a human editor, and back. The diagram on the next page shows how we anonymize private and sensitive data before translating content. The next section will explain this in more detail.

OUR TRANSLATION WORKFLOW

When we receive a customer message, the first thing we do is to automatically remove key categories of personal data. This means that those are not seen by a human until it reaches its intended destination.This includes the training data we use for our AI translation engines.

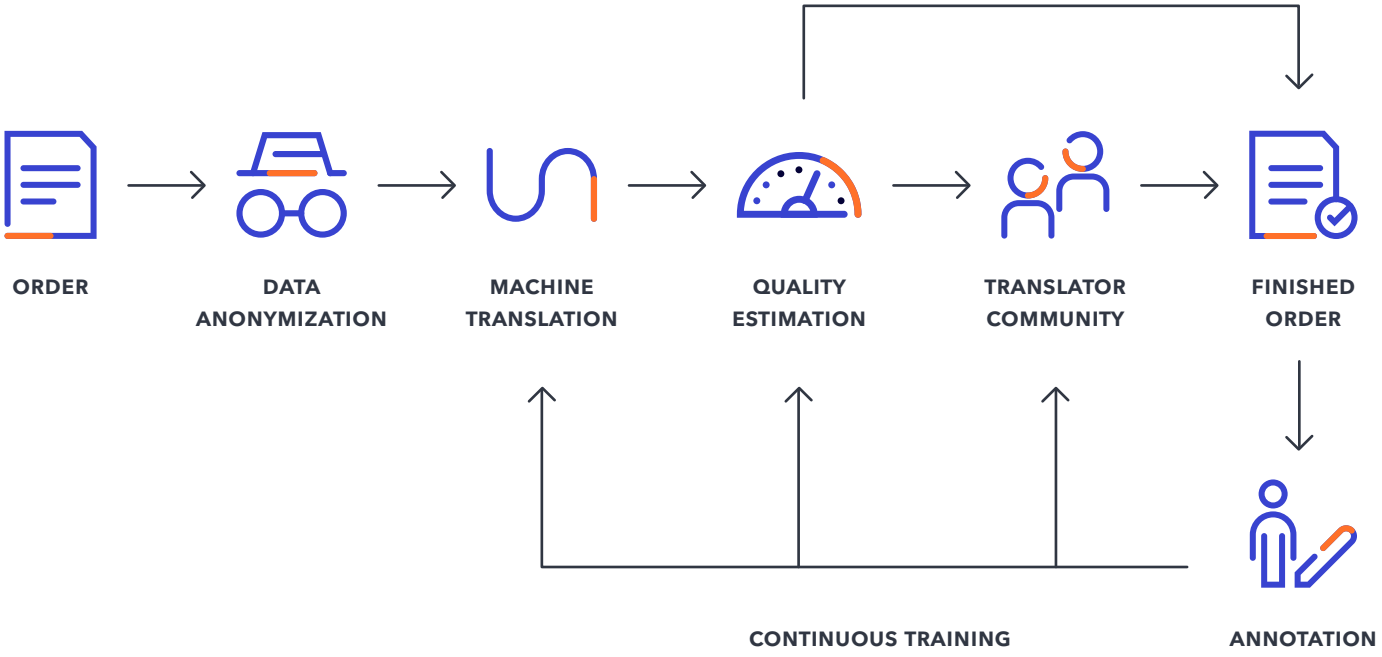

Customer Data Security

Here is how we enforce customer data security at Unbabel.

- **Pseudonymization:** All content passing through Unbabel's translation pipeline from its customers goes through an automated pseudonymization process which removes sensitive data and restores it before delivery (see diagram above). No sensitive data captured by our pseudoanonymization tool is shared with the human editor community. We adhere to EU GDPR and CCPA to protect personally identifiable information (PII).

  Eraser is the tool Unbabel currently uses that automatically hides sensitive data from our editor community, including email addresses, phone numbers, and credit card information. It redacts specified classes of PII, from messages before they enter our translation pipeline.

- **Access control:** All access to Unbabel's products and services is encrypted and protected by a firewall. All access credentials are segregated by work-group areas, provided to staff on a need-to-know basis, and audited based on internal security heuristics.

- **Two-factor authentication:** Access to administration applications is secured by 2FA on top of standard user account authentication.

- **Audits and external validation:** Unbabel applies internal security policies to increase penetration barriers, digital and physical, and regularly performs information security audits using third-party vendors to validate our compliance with best practices and relevant regulations.

- **Encryption:** Data is encrypted in transit and at rest. Much of the data that we translate contains personally identifiable information (PII), and we recognize our responsibilities as a data processor under GDPR. To ensure that we meet these responsibilities, we secure all customer data as if it contained PII, using industry-standard and up-to-date cryptographic mechanisms.

DATA RETENTION

We retain data to train our machine translation engines for as long as we provide services to a specific customer. We also retain some data for research and audit purposes and to refine our machine learning algorithms. At the request of a customer, we can securely remove any or all data from our systems. Unbabel complies with the principle of data minimization.

THE UNBABEL EDITOR COMMUNITY

Great translation can't come at the cost of security. Unbabel uses human editors to support machine translation in driving quality for its service. Not every message goes to an editor to translate. Editors translate messages only when our AI translation engines detect that the machine translation may be sub-par. Editors are also used to grade the work of machine translation and provide annotations that improve the service.

The data that Unbabel translates is privileged, and we respect the relationship that our customers have with their customers. Unbabel treats customer data as confidential, following industry best practices and relevant regulations at all times. In working with editors, Unbabel uses a layered set of controls to reduce the security risk to the data.

Here is some more detail on how we train, equip, and monitor our editor community to uphold Unbabel's commitment to security and privacy.

ONBOARDING AND OFFBOARDING

The editors that work for Unbabel are required to sign NDAs before accessing Unbabel systems. The NDA and associated contracts protect our customer data and provide Unbabel with the necessary legal authority to remove editors who do not follow policies. Copies of the NDA and the latest terms that all editors sign can be viewed by emailing **SECURITY@UNBABEL.COM**

Before an editor can start work on the Unbabel platform, they must complete training on the responsibilities of an editor and the rules for working on the platform. This includes detailed information regarding their responsibility to protect all communications they are asked to translate.

Editors are tested on their ability to translate a series of texts that mimic the challenges they'll face in customer exchanges. Their work is then evaluated by an elite team of linguists according to Unbabel's guidelines and quality standards. Editors are only able to work on customer jobs if they match or exceed our threshold score. Working editors receive periodic evaluations from our linguists, who will provide another round of feedback on tasks. If they no longer meet minimum requirements for that language pair, Unbabel deactivates access to paid tasks.

ONGOING MONITORING

All editor activity on the Unbabel systems is recorded and audited. Unbabel monitors all systems that the editors use for suspicious behaviors. Any suspected infringement or failure to meet the minimum requirements for working on the Unbabel platform will result in that editor being deactivated from access to information. The community team at Unbabel regularly reviews all monitoring systems in place to ensure that it provides the protection that our customers require.

TRANSACTIONAL SECURITY

As mentioned, Unbabel has built a service called Eraser that removes personally identifiable information and other sensitive data before machine translation and before being sent to an editor for translation. The only thing editors will ever see is a token, which is a placeholder text that identifies the type of content customers are describing. This ensures translators have the context they need to work while preserving privacy and security.

Editors work in a secure interface where they:

- Accept the job
- Translate the fragment of the message that constitutes the job
- Submit the job back to Unbabel

In our contract (and reinforced in editor training), Unbabel requires that all work must be conducted on the Unbabel platform. Editors cannot download or upload their tasks to/from other services.

Network Security

A minimum-required access policy is enforced throughout systems, using firewall ingress control rules to limit access on a per-IP/port basis. All of our systems are firewalled by our cloud providers. We use AWS VPC (Security Groups, Subnets, ACLs) secure configurations and Amazon GuardDuty for host protection. We use CloudFlare for DDoS defense, as well as hardware and software load balancers on AWS.

Unbabel uses TLS-negotiated secure sockets for data transfer, enforcing modern cryptographic configurations. We have disabled old versions of the TLS protocol (SSL v2, SSL v3, TLS v1, etc.) and insecure ciphers. Unbabel uses AES-256 for data encryption at rest, via Amazon EBS transparent encryption. Columnar encryption is enforced for personally identifiable information (PII) flagged by our Eraser service.

Our network is not accessible from the internet on insecure or clear-text protocols such as telnet, snmp, smb/cifs, etc.

We employ intrusion detection systems, data loss protection (DLP) solutions, and network logging for extra layers of security.

Application Security

At Unbabel, we observe and enforce secure application development practices. Unbabel follows Conversational Development best practices, which means:

- Updates are tested locally and on staging before being deployed on production. (The source code is managed on Gitlab.)
- The rollback procedure can be achieved by deploying the last working deploy.
- We have coding style guides and mandatory code reviews. We run multiple checkers and validations on our software build pipelines.
- Code is reviewed using static code analysis tools.
- Code is reviewed by a different team member than the author to approve before being merged.

We have several test phases (Unit, Functional, Performance). We also conduct System Integration Tests and User Acceptance Tests before deploying into production.

For security measures against common attacks, we use different tools, including Cloudflare Flan and OWASP Zap, that scan for OWASP Top 10 Web Application Security Risks and enforces security best practices.

Additionally, we employ the following safeguards:

- Secure logging of confidential and restricted information
- File transfer encryption
- Endpoint encryption
- Database encryption
- Credential security
- Permissions logging
- Multifactor security

Unbabel regularly applies system-wide patches and runs security tests on all systems. Continuous monitoring and tracking are mandatory in all systems, as well as log-based anomaly detection. Change control procedures are embedded and mandatory in our CI/CD development pipeline.

Unbabel regularly applies system-wide patches and runs security tests on all systems. Continuous monitoring and tracking are mandatory in all systems, as well as log-based anomaly detection. Change control procedures are embedded and mandatory in our CI/CD development pipeline.

Vendor Security

To support the delivery of our services, Unbabel relies on service providers. Any service provider engaged by Unbabel that might have access to or process data that may contain personal data is considered a Sub-Processor.

The Unbabel translation pipeline is designed with robust privacy and security measures. That said, Unbabel still performs a security and privacy review of the practices of any Processors before engaging with them. Our supplier security policy governs our vendor management and imposes Unbabel's security requirements in order to minimize risk to our customers' data. A full list of current processors can be found in our Privacy Policy.

Any processor or subprocessor used by Unbabel is put under rigorous scrutiny to assess their security and privacy policies, as well as the adoption of adequate safeguards. We require all of our Processors to have signed a data processing agreement (DPA) with us, similar to the DPA that our customers sign with us. They are required by this DPA to:

- Process personal data as defined both in the DPA and the applicable laws
- Restrict data access only to trusted and contractually bound staff to assure data privacy and security
- Train the staff who has access to personal data on data privacy and protection best practices and policies
- Implement processes that take privacy into account throughout all their data processing activities
- Inform Unbabel about any actual or potential data breach
- Cooperate with Data Protection Authorities or Data Controllers when required

Unbabel only discloses personal data to service providers where the disclosure is absolutely necessary to provide the services that our customers request. Unbabel will not sell any kind of personal data.

Compliance & Certifications

Unbabel is fully compliant with GDPR, and certified under the EU-US Privacy Shield and ISO 27001. Unbabel also complies with the California Consumer Protection of 2018 (“CCPA”). PCI DSS compliance is in progress at the time of publication.

All Unbabel data is processed using Google Cloud and AWS on servers based in the US and Ireland.

For additional relevant information on security and privacy, please see our **PRIVACY POLICY** or contact us at **SECURITY@UNBABEL.COM**