Where's My Data and How is it Stored and Protected?

ISMS - 025

enate

Contents

1. Platform Design

This document aims to answer fundamental questions asked by many Enate customers when they are considering the Enate SaaS service. It should be read in conjunction with our ISO27001 accreditation and associated information security policies. By reading this document you should be able to understand:

- Where your data is stored.
- How it is separated from other Enate customer's data.
- How it is protected and kept available to meet our SLA promises to you.
- How Enate uses encryption and cryptography to secure your data.
- How we meet our obligations under GDPR and other data protection regulations.

The diagram below shows the overall picture. The remainder of the document describes this in more detail.

1.1 Enate Hight Level Architecture Diagram

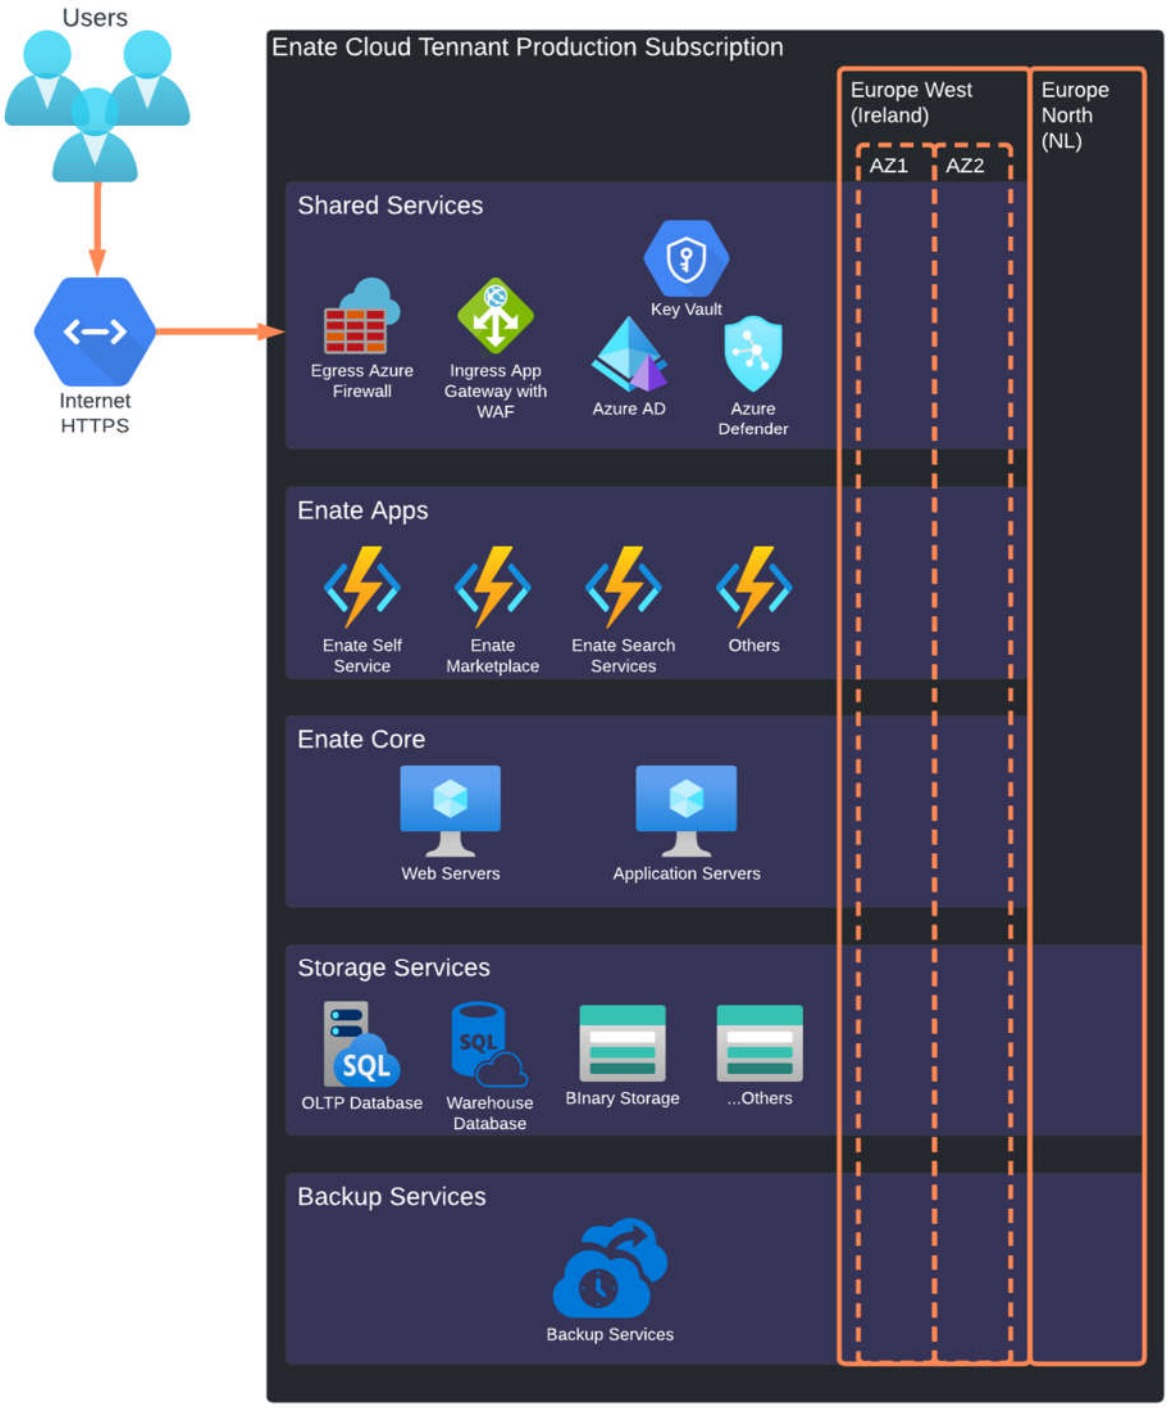

2. Where is my data stored?

Your data is stored in a number of locations and in different conditions. This section explains where and when.

2.1 Regions & Availability Zones

Your data is held in two regions, Ireland and the Netherlands

- Production – the Enate production services are operated between the Azure Europe West (Ireland) region and the Azure Europe North (Netherlands) region. The workloads operate in Europe West by default with Europe North available for DR scenarios.
- Availability Zones – unless otherwise stated in this document, all Enate services run in at least two availability zones (i.e. two datacentres) providing redundancy against failure of a single availability zone.
- Disaster Recovery (DR) – Backups & copies of Enate production data are stored in both the Europe West and Europe North regions.

For more information on the data centres storing your data please see these links:

More information:
Azure

- Global Infrastructure | Microsoft Azure
- Physical security of Azure datacentres - Microsoft Azure | Microsoft Docs
- Azure compliance documentation | Microsoft Docs

2.2 How is Data Stored in Enate?

Enate has two tiers of data storage in operation. The Primary Data Stores hold the master copy of customer data. Secondary Data Stores hold subsets and transformed copies of customer data in order to support specific features of the Enate platform. All secondary data stores can be recreated from the primary data stores.

The primary data stores are:

- The Enate Online Transaction Processing (OLTP) Database – Used for storing and managing all work and configuration data.
- Azure Blob Storage – Used for managing and storing files, emails and large format data.

There are many secondary data stores, some examples include:

- Enate Warehouse Database – A simplified and transformed version of OLTP used for reporting and MI.
- Power BI Datasets – Used to hold copies of OLTP data to drive reports against.
- Azure Cognitive Search Indexes – Used to support sophisticated free text searching across the platform.
- Azure Cosmos DB – Used for managing searchable data in NoSQL format.

2.3 Principles Applied to All Data Stores

There are fundamental principles applied to all Primary and Secondary data stores used in the Enate service.

- Encryption at Rest – Regardless of the data store, all customer data is encrypted at rest (see section 5 for a discussion of cryptographic approaches).
- Data Segregation – Each Enate customer's data is logically segregated from other customers' data using the most appropriate segregation approach for the storage technology being used. Your data is NOT merged into a single store with other customers. For example:
    - OLTP Database - Each Enate customer has a separate Azure SQL Database.
    - Power BI Datasets – Each Enate customer has a separate Power BI Embedded Workspace in which their datasets reside.

2.4 Primary Data Store - Enate OLTP Database

The master copy of all live customer data (apart from files and emails) is in the Enate OLTP Database. In line with the segregation approach outlined above, there is a separate database for each instance of Enate so your data is NOT merged into a common database with other customer data.

The Enate OLTP database is a Microsoft Azure SQL Database, this is Microsoft's cloud native database technology. Depending on the scale of your Enate platfrom your Enate OLTP database will either be managed as:

General Purpose Elastic Pool: configured with Zone Redundant Availability and Regional backups. You can find out more about this configuration here High availability - Azure SQL Database and SQL Managed Instance | Microsoft Learn … but a simple translation means:
General Purpose – The database uses the standard performance tier for Azure SQL Database.

- Elastic Pool – The database is running on the same server as some other Enate customer's databases.
- Zone Redundant Availability – Means that the database is High Availability and will automatically switch to another Availability Zone (i.e. datacentre) if there is a problem in the zone it is currently operating in. This switch is near instant.
- Regional Backup – means that backups are stored in multiple Availability Zones in both Ireland and in the Netherlands region. Effectively meaning that there are at least 4 copies of all backups.

- *Hyperscale*: configured with Zone Redundant Availability and Regional Backups. Hyperscale is the Azure SQL Database configuration for large databases requiring high performance as they scale. You can read more about the Azure SQL Database Hyperscale configuration here Hyperscale distributed functions architecture - Azure SQL Database | Microsoft Learn

All databases are encrypted using SQL Server Transparent Data using AES256 encryption. Key management is operated with the Azure Key Vault service you can read more about this service here: What is Azure Key Vault? | Microsoft Learn.

A. Backup Policy

Enate makes use of the backup feature of Azure SQL Database. You can learn more about these features here. https://docs.microsoft.com/en-us/azure/sql-database/sql-database-automated-backups .

Currently the service is configured to take Full Backups every week, Differential Backups every day and Log Backups every 5 minutes. Meaning a worst case RPO of T-5 minutes. Microsoft may further improve these parameters as they enhance the service.

Point in time restore can be achieved to any point in the last 35 days.

2.5 Primary Data Store – File & Attachment Storage

Enate stores encrypted files and emails in external binary storage databases outside the OLTP database. Files and emails are encrypted AES256 using a symmetric key that is managed in the Azure Key Vault service.

Enate automatically stores all files and emails using the Microsoft Azure Blob Storage service with replication across three Availability Zones. Each file is written twice (not replicated) to the Primary Region (Europe West) and the Secondary Region (Europe North).

In line with the segregation policies set out above, each Enate customer's data is stored in a separate 'Bucket' (yes … this is the technical term) within the storage services.

Using this write twice approach provides additional protection against the unlikely event of data corruption, allowing data corrupted in one region to be recovered from the other region.

For UAT environments, a read-only connection to the Microsoft Azure Storage can be configured to assist with testing upon written request from a customer, however the default is to store only test data in a UAT database.

2.6 Disaster Recovery

A. Primary Data Stores (DR)

Disaster recovery is only enacted if an entire Microsoft Azure region fails. All Enate services are designed to cope seamlessly with failure of a single Availability Zone within a region.

OLTP Database

Enate makes use of the Active/ Passive configuration in Azure SQL Database. In the extremely unlikely event of the loss of the Europe West region, the service will automatically be failed over to the Europe North region.

You can read more about Azure SQL Database Point In Time Restore here Azure SQL Database Point in Time Restore | Azure Blog and Updates | Microsoft Azure

Email and File Storage

In the event that the Enate platform cannot read binary data form the Azure Storage in the Europe West region, it will automatically read the file from the Europe North region. Thus DR is instantaneous and achieved on a file-by-file basis.

B. Secondary Data Stores

The DR approach to secondary data stores can be flexible because ALL secondary data stores can always be automatically re-created from data in the primary data stores. The approach therefore depends on the specific Azure service used to support the Enate feature. If the underlying Azure service offers Region Redundant Backup then Enate makes use of this feature for DR purposes. If the underlying Azure service does not, then the approach to DR is to automatically re-create the Secondary data store from the data held in the primary data stores.

C. Compute

In the event that all Availability Zones in the Azure West region are unavailable, compute will is provisioned in the Azure North region.

D. RTO and RPO

Recover Point Objective (RPO) is 5 minutes. i.e. No more than 5 minutes of data should ever be lost.

Recover Time Objective is 4 hours i.e. in the event of a complete disaster losing a complete Azure region the service should be restored within 4 hours.

3. How is my Data Accessed?

Data is only accessed through the Enate Application or the Enate API Surface.

3.1 Application Access

A. Enate Work Manager and Builder

The main route for day-to-day access is through the Enate Work Manager or Enate Builder websites. All access to these sites is authenticated (either with Username and Password or Single Sign-on) and the information that a user sees is strictly controlled by the permissions and role granted to their account. All communications to and from the website are encrypted over TLS 1.2 if supported by the client connecting and TLS 1.1 minimum.

The ONLY users who can access your data through the application are people who you have chosen to grant user accounts and permissions to.

B. Enate Self Service

Enate Self Service is an application designed to allow your customers to start, track and update service requests with you. All access to Enate Self Service is authenticated using the OAuth2.0 standard through Azure Active Directory B2C and is only supported with Single Sign-on. Self Service users are only able to see and interact with data that they have been explicitly tagged on in the Enate application. All communications to and from the website are encrypted over TLS 1.2 if supported and TLS 1.1 minimum.

The ONLY people who can access your data through Self Service are people who you have chosen to grant user accounts and permissions to.

3.2 Enate API Surface

Enate supports three APIs and which can be used to access data in different ways.

A. Enate Core API

This is the API used to interact with the deep features in Enate Work Manager and Enate Builder. The API includes both inbound calls and webhooks.

This is the API that is consumed by the Enate Work Manager and Enate Builder applications. Access to the API is authenticated either in the same way as the Work Manager or Builder application or by user granted API Key. Access granted by API Key assumes the same permissions and role as the User account that the API key is generated for.

The ONLY users who can access your data through the application are those who you have chosen to grant accounts and permissions to. All communications over this API encrypted over TLS 1.2 if supported and TLS 1.1 minimum.

B. Enate Self Service API

This is a simple API used to support the Enate Self Service application or embed access to Enate self-service within other applications. The Self Service API is authenticated by OAuth 2.0 with each account only having access to data that they have been explicitly tagged on.

C. Enate Marketplace API

This is the API used by Digital Workers such as IDP, RPA, Document embellishment to interact with Enate. The API implements Enate Patterns that are strongly versioned and designed to be called by an Adapter as the interface between Enate and the 3rd party digital worker.

This API is only accessible when adapters are registered with the Enate Marketplace and reviewed by the Enate team, at which point a certificate is exchanged. All calls between adapters and the marketplace are authenticated by Mutual TLS 1.2.

The Marketplace API ONLY makes data available to adapters if you have explicitly enabled a specific Enate Pattern in the Enate Builder application.

3.3 Enterprise Shared Data Access

For Enterprise tier customers two additional methods of data access are possible:

- Power BI X Tenant Dataset Sharing – This is an approach whereby specific users in your Active Directory can be granted guest access through the enate.cloud active directory to a dataset in your Enate Power BI Workspace. This makes the dataset queryable by these users within your Power BI Tenant. You can find more information on this service here.
- Azure Data Share – This is an approach whereby you can have a near real time synchronised copy of your Enate Data Warehouse in your Azure Tenant. You can find more information on this service here What is Azure Data Share? | Microsoft Learn

Both of these data access methods are secured through explicit arrangement with you and neither is enabled by default within the Enate service.

3.4 Enate Administrator Access

Enate has a small team of infrastructure support administrators who have administration rights to the Enate SaaS infrastructure, servers and databases. These staff go through rigorous security checks before joining the Enate team (please see our policies for more information).

These users do not have routine access to your data. Access is under management supervision. They will only ever access your data in order to resolve tickets or incidents raised by you and then only if access to the underlying data is essential to resolve the issue. All access to the infrastructure by this team is automatically logged and audited.

Enate Administrators access the system through a bastion on the Enate Production platform from where they can administer the servers.

A. Enate Monitor desktop application

Enate Monitor is a tool used by some members of the Enate support team. All access to Enate Monitor is controlled through username and password authentication and the information that a user sees within Enate Monitor is strictly controlled by the permissions granted to their account. Enate Monitor accesses data over a web service which is encrypted over TLS 1.2.

4. Security and Posture

1.2 Security Information and Event Management

Log aggregation provides real-time analysis of events generated by the Enate system, underlying operating system and Azure audit trails for access logging and trend analysis. A combination of rules-based and machine learning logic detects potential threats and anomalies and raises these to operational teams for further investigation. Enate uses Datadog and Azure AppInsights for log management and analysis.

1.3 Anti-malware and Threat Detection

Azure Defender for Endpoint provides next-generation OS-level threat detection with behaviour-based, heuristic, and real-time antivirus protection. The cloud-delivered protection provides near-instant detection on new and emerging threats. Signature updates to complement these capabilities are deployed daily. You can read more about Aure Defender for Endpoint here Using Microsoft Defender for Endpoint in Microsoft Defender for Cloud to protect native, on-premises, and AWS machines. | Microsoft Learn.

1.4 Web Application Firewall

Enate uses the Azure Web Application Firewall (WAF) to protect against common web hacking techniques. This provides protection for the Open Web Application Security (OWASP) top 10 security issues, including cross-site scripting and SQL injection attacks. Additionally, custom rules may be deployed to handle specific emerging threats faster than a fix may be able to be deployed by a vendor. You can read more about Azure WAF here Azure Web Application Firewall (WAF) | Microsoft Azure.

1.5 Egress Traffic and Intrusion Detection and Prevention

Egress traffic within the Enate environment is centralised and subject to network-based detection and intrusion prevention scanning (IDPS) including the capability to inspect TLS encrypted traffic. Additionally, traffic to specific categories of website or to likely compromised/malicious domains and URLs is restricted via URL filtering. Enate uses the Azure Firewall Premium service to provide these capabilities. You can read more about Azure Firewall Premium here What is Azure Firewall? | Microsoft Learn

1.6 Cloud Security Posture Management

Enate utilises Azure Defender for Cloud to provide Cloud Security Posture Management (CSPM). This provides continuous scanning against a wide range of industry standard benchmarks and regulatory standards and provides security recommendations to fix any misconfigurations or weaknesses. You can read more about Azure Defender for Cloud here What is Microsoft Defender for Cloud? - Microsoft Defender for Cloud | Microsoft Learn.

4.1 Cloud Security and Image Hardening

Deployments within the Enate environment are hardened according to industry standard benchmarks provided by the NIST 800-53 and implemented through the Defence Information Systems Agency STIGs (Security Technical Implementation Guides). For cloud platform security, these are verified continuously via Azure Defender for Cloud CSPM. For image-based deployments, these are verified during the image

creation process with additional verification after deployment via Azure Defender for Cloud Vulnerability Management and Qualys Vulnerability Assessment.

4.2 Auditing Within the Enate Application

Within the application, Enate provide fine grained work item access auditing capability keeping track of which users have accessed which work items when.

5. Authentication and Cryptographic Standards

5.1 User Authentication

Enate supports two authentication regimens:

- Application Authentication – User accounts are held in the Enate application. Passwords are hashed using SHA2 with randomly generated salt. Password policies are configurable and include minimum password length, combination of case and special characters, expiry and minimum number of passwords before one can be re-used.
- SAML 2.0 Single Sign-on Authentication – The authentication of users is managed through SAML 2.0 and user passwords are not stored in Enate.

The platform supports both authentication regimens in parallel although SAML only is a configuration option for the Enate Work Manager site. Application Authentication is required to be used for access to the deprecated desktop client applications.

5.2 Cryptographic Standards

Passwords and keys are always either encrypted or hashed when stored in or by the Enate application.

A. Encryption

Where encryption is deployed the algorithm used is AES-256 with randomly generated initialisation vectors.

B. Hashing

Where Hashing is deployed the algorithm used in SHA-512 with a randomly generated salt.

C. Encryption of Data at Rest

Data at rest is encrypted using the AES-256 algorithm with randomly generated key initialisation vectors and keys managed through the Azure Key Vault service with full auditing of key usage.

D. Key Management

All cryptographic keys are managed and stored in Azure Keyvault.

E. Encryption in Transit

Data is encrypted in transit using TLS1.2

5.3 Administrative Authentication

Enate Administrators use 2-factor authentications for all administration of the Enate platform.

6. General Data Protection Regulation ('GDPR')

6.1 GDPR

The General Data Protection Regulation (EU) 2016/679 – commonly known as 'GDPR' - applies in member states of the European Union, member states of the European Economic Area, and in some circumstances where the Regulation affects organisations outside the European Union because they process data relating to individuals entitled to protection by the Regulation. The United Kingdom has left the European Union but has retained European data protection legislation using the UK's Data Protection Act 2018 and what has become known as 'UK GDPR'. Enate serves customers who need to be compliant and Enate is committed to maintain full compliance with GDPR.

The Enate Solution is commonly used by customers who operate in multiple jurisdictions. This means that some personal data might be within the scope of protection provided by Data Protection Legislation and some might not. This might change the text in your Enate contract terms but it does not change Enate's operational procedures: Enate adopts the same procedures for all customer data, regardless of whether it happens to be protected by Data Protection Legislation.

When Enate provides solutions to customers, we act as a data processor. The Customer is the data controller. Processors and Controllers each have obligations under GDPR.

Customers assessing the application of GDPR to use of an Enate solution should keep in mind that the solution is generally agnostic about the types of data that can be processed. Other than user credentials, which are plainly 'personal data', Enate (the company and the solution) does not inherently recognise what data is 'personal data'. For this reason, Enate processes all customer data on the assumption that it might include personal data. Enate will defer to customers any requests by data subjects to exercise rights such as access requests. Similarly, Enate does not make any distinction to specifically recognise any of the so-called 'special categories' of data listed in Article 9 of GDPR.

GDPR requires transparency when a processor appoints another processor. This document describes how the hosted solution uses facilities provided by Microsoft. Our formal disclosure of such details required by Article 28(2) of GDPR is maintained here: https://docs.enate.net/enate-contract-documents/

A. Integrity and Confidentiality

Data should be processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures ('integrity and confidentiality')

The purpose of the Enate ISMS and the design of the Enate platform support this principle. Product features such as end-to-end encryption, single sign-on, user authentication and role-based permissions support this principle. Enate's Information Security Management System validated under ISO27001 is also a fundamental part of

supporting this principle providing setting out both technical and organisation controls to prevent inappropriate processing of data.

B. Storage Limitation

Data should be kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed; personal data may be stored for longer periods insofar as the personal data will be processed solely for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes subject to implementation of the appropriate technical and organisational measures required by the GDPR in order to safeguard the rights and freedoms of individuals ('storage limitation')

The Enate platform includes features to allow transactions in the Enate system to be tagged as containing personal data and the person to whom the data relates. It is customer's responsibility to use such features to ensure compliance with the GDPR. By default all data remains within Enate for the duration of the contract. At the end of the contract the data is either handed to the customer in machine readable form or deleted, based upon the customer's preference. It is possible to purge data from Enate at other times, however this is not an online operation within the system and is dealt with on a case by case basis. In the rare case that we see such a request, we require the customer to be specific with regards to the data they would want removed as we do not distinguish for example between personal data and other types of data within the system, other than for user credentials. This process would undergo specific development and testing ahead of being executed against the production database.

C. Accuracy

Data should be accurate and, where necessary, kept up to date; every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay ('accuracy');

Enate includes features to allow data to be corrected and amended as appropriate by the Data Controller. Audit trails stored within the system however cannot be amended.

D. Individual Rights

The GDPR sets out individual rights to be informed, right to access, right to rectification, right to erasure.

Enate supports these rights through software features allowing clients to meet these requirements without recourse to Enate. Clients must take care to understand how Enate will fit into the wider processing of data subject access requests within their organisation.

Right to Erasure and Rectification – Enate does not provide features to erase or rectify personal data from database backups. Erasure and rectification propagate through the backup cycle.

Please note that most Enate customers use Enate as a transaction processing system in support of commercially agreed (whether internally or externally) transactions and as such personal data stored within these transaction records may be required for the

exercise or defence of legal claims and therefore may be exempt from the right to deletion.

E. Other Principles

The other principles of the GDPR namely:

- Lawfulness, fairness and transparency
- Purpose limitation
- Data Minimisation

Remain customer's responsibility as Enate is providing the facilities for processing data in line with the decisions that customers make with respect to these principles.

F. Data Transfers

Data is processed in the UK and the EEA apart from technical support services provided by our wholly-owned subsidiary Enate Technologies India Private Limited. Enate Technologies India Private Limited will sign GDPR-recognised standard contractual clauses with customers if needed.

6.2 Other Data Protection Regimes

Enate is not an expert in other data protection regimens from where customers may be submitting data to Enate solutions. Customers with other local regulatory requirements must independently assess whether Enate' solution is suitable but Enate will provide information required to enable that assessment.

7. Outbound Email Routing

It is often most efficient to send email outbound direct from Enate rather than through your email systems. This is because Enate often becomes the system of record for most Email when deployed. If you choose to send email directly from Enate rather than through your mail servers, then the service will be configured as follows:

Enate makes use of the SendGrid service to deliver emails securely.

To enable this service and ensure that emails are sent properly, changes need to be made to your SPF and DKIM records. These changes ensure that emails do not trigger spam filters. The changes required are as follows:

- Add one CNAME record for validation (em1234.customer.com)
- Add 2 CNAME records for DKIM (s1._domainkey.customer.com, s2._domainkey.customer.com)

Using this approach simplifies setup; you never need to update an SPF record specifically for this because the service looks back at the validation and works it out, and you don't need to manage and rotate keys for DKIM.

SendGrid does not persistently store emails – they are stored only for the time it takes to send them. Emails are sent encrypted in transit over the SMTPS protocol. All email content is encrypted at rest while queued for sending.

8. EnateAI – Using Azure OpenAI Service

For Customers who are making use of the EnateAI capabilities in the Enate Marketplace data within these services Is managed as follows:

8.1 Background

EnateAI makes use of the Azure OpenAI service under the hood. The Azure OpenAI service allows Enate to make use of private Instances of the OpenAI models. This Is managed as a secondary data source as set out in the rest of this document.

You can read more details about the security of this service at Data, privacy, and security for Azure OpenAI Service - Azure Cognitive Services | Microsoft Learn .

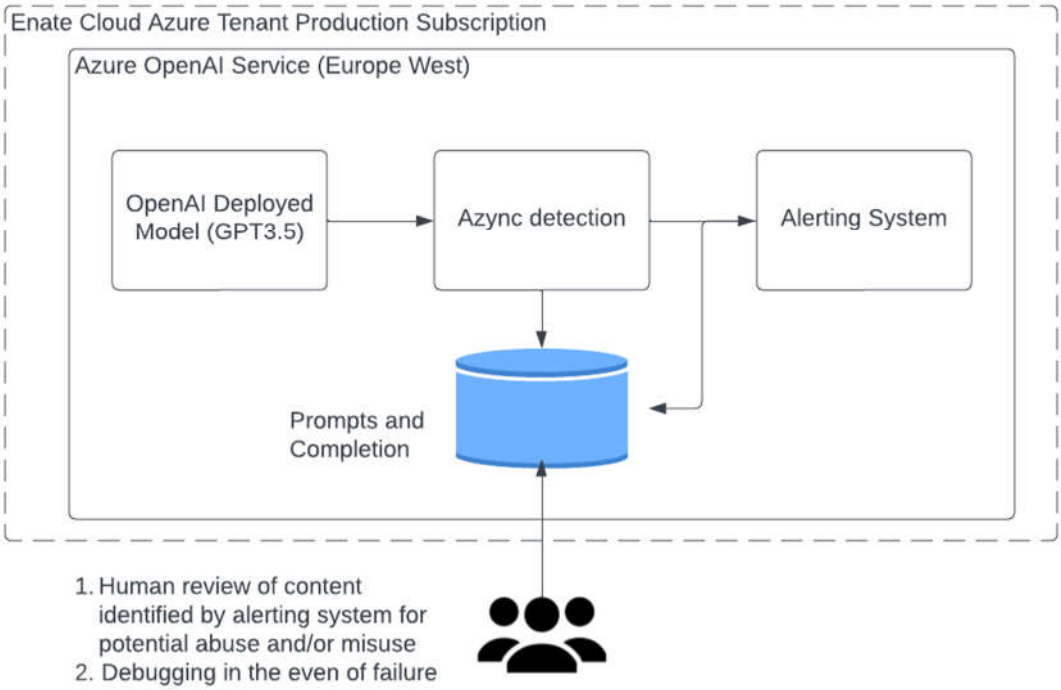

When you're reviewing this document, you will see that there are other services available In the Azure OpenAI service namely Training Data Upload, Training Infrastructure and Custom Model. Enate does NOT make use of these services, your data Is not used or stored to train dedicated versions of the OpenAI model.

9. Enate SaaS Options

Outside of our standard SaaS offering described above, we have three additional options available to customers. These options are chargeable and will therefore require more specific discussion relating to the architectural requirements of individual customers. They are:

- Bring your own key (BYOK)
- Private Azure Binary Storage Account – Bring Your Own Bucket (BYOB)
- Enate Private Cloud

9.1 Bring Your Own Key (BYOK)

Enate supports customers Cross Tenant 'Bring Your Own Key' for encryption of the Enate OLTP database and secondary data stores that support this feature.

To make use of this capability customers will need to provide and manage encryption keys using Azure KeyVault in their Azure tenant. This key is then made available to Enate and used to encrypt the Enate databases.

Should a customer revoke access to that key the data becomes instantly unreadable in the Enate platform and the platform will become inoperable.

You can read more about how Cross Tenant Customer Managed Keys works here Cross-tenant customer-managed keys with transparent data encryption - Azure SQL Database & Azure Synapse Analytics | Microsoft Learn.

NOTE: If your organisation is not highly proficient at key management you should NOT adopt this option. Loss or corruption of the key will result in immediate and irrevocable loss of ALL of your data.

9.2 Bring Your Own Bucket (BYOB) - Private Azure Binary Storage Accounts

Enate is always provisioned with the primary binary storage configured in the Enate Azure tenant as described in 2.5.However, should you wish this storage to be in your Azure tenant Enate can support this.

To enable this feature, you will need to:

- Create two Azure Storage Accounts in two separate Azure Regions within your Azure tenant. We recommend that one of these regions is Europe West to maximise performance.
- Create an Azure App Registration that is granted access to these storage accounts.
- Configure Enate to use these storage accounts rather than the Enate Default.

NOTE: If your organisation is not proficient at managing Azure storage then you should NOT adopt this option. Deletion or corruption of data in these storage accounts will result in immediate and irrevocable data loss.

9.3 Enate Private Cloud

Enate private cloud is a chargeable service option for the most security conscious clients. It completely segregates the compute and data storage tiers for an individual Enate Tennant. The only components that remain shared are the Enate Apps (of which there is only ever one code base deployed and egress through the firewall. The diagram below shows the high-level architecture. Enate Private Cloud supports BYOK and BYOB in addition.