LetsBe Biz — Security and GDPR Framework

Security, Privacy and Compliance

Version: v1.1
Date: February 26, 2026
Company: LetsBe Solutions LLC
Contact: matt@letsbe.solutions

221 North Broad Street, Suite 3A, Middletown, DE 19709

Contents

1. LetsBe Biz — Security & GDPR Compliance Framework

Version: 1.1 **Date:** February 26, 2026 **Authors:** Matt (Founder), Claude (Architecture) **Status:** Living Document **Companion docs:** Technical Architecture v1.2, Foundation Document v1.0, Pricing Model v2.2

1.1 1. Purpose & Scope

This document defines the security posture, data protection obligations, and regulatory compliance framework for the LetsBe Biz platform. It covers:

- The security architecture enforced on every tenant VPS
- GDPR obligations and how the platform satisfies each one
- Data flows, retention, and deletion capabilities
- Subprocessor management (LLM providers, hosting, payment)
- EU AI Act readiness and transparency requirements
- North American privacy compliance (CCPA/CPRA and state laws)
- Customer-facing artifacts: DPA template, privacy policy inputs, security FAQ

Target audiences: internal engineering (for implementation guidance), legal counsel (for DPA and policy drafting), sales (for customer security questions), and customers themselves (via the published security page).

Regulatory scope: The platform serves SMBs primarily in the EU and North America. The compliance framework is designed for GDPR (EU/EEA), the EU AI Act, CCPA/CPRA (California), and the growing patchwork of US state privacy laws (Virginia CDPA, Colorado CPA, Connecticut CTDPA, Indiana, Kentucky, Rhode Island, and others effective through 2026). Where requirements diverge, the stricter standard applies.

1.2 2. Data Architecture & Classification

1.2.1 2.1 Data Categories

LetsBe processes several distinct categories of data. Each category has different handling rules, retention periods, and legal bases.

Category	Examples	Where Stored	Controller	Processor
Customer Account Data	Name, email, billing address, payment method	Hub (central platform)	LetsBe	Stripe (payment)

Category	Examples	Where Stored	Controller	Processor
Business Profile Data	Business name, industry, team size, bio	Hub + tenant VPS	Customer	LetsBe
Tool Data	CRM contacts, emails, calendar events, files, invoices, wiki pages	Tenant VPS only	Customer	LetsBe
AI Conversation Data	Chat messages, agent responses, session transcripts	Tenant VPS only	Customer	LetsBe
AI Reasoning Data	LLM prompts sent to external providers (redacted)	Transit only — not stored by LetsBe	Customer	LetsBe → LLM provider (subprocessor)
Credential Data	Tool passwords, API keys, OAuth tokens	Tenant VPS only (encrypted SQLite)	Customer	LetsBe
Telemetry Data	Token usage, agent activity counts, error rates	Hub (aggregated, no PII)	LetsBe	—
Server Infrastructure Data	IP addresses, SSH keys, nginx configs, Docker state	Tenant VPS only	LetsBe	Netcup (hosting)

1.2.2 2.2 Data Residency

All tenant data stays on the customer's VPS. This is not a policy choice — it's an architectural decision (Decision #18: one customer = one VPS, permanently). The VPS is provisioned at the hosting provider's data center.

Data Type	Location	Jurisdiction
Tenant VPS (all tool data, AI conversations, credentials)	Customer's choice: Netcup Germany/Austria (EU) or Netcup Manassas, Virginia (US)	EU or US — depends on customer region selection
Hub (account data, billing, telemetry)	EU-based hosting (Germany)	EU
Payment processing	Stripe	EU entity for EU customers, US entity for NA customers
LLM inference (redacted prompts only)	Varies by provider (OpenRouter → Anthropic, Google, DeepSeek, etc.)	Mixed — see §5 Subprocessors

Key privacy advantage: Each customer gets their own isolated VPS in the data center region they choose at signup. EU customers (Netcup Germany/Austria) benefit from native GDPR jurisdiction; North American customers (Netcup Manassas, Virginia) get low-latency access with US-jurisdiction hosting. In both cases, the only data that crosses borders is redacted LLM prompts (with all secrets, credentials, and configurable PII categories stripped before transmission). EU customers who need strict data residency should choose the EU region; NA customers who prefer GDPR protections can also opt into the EU region.

1.2.3 2.3 Data Flow Diagram

```
Customer (browser/app)


Hub (EU - Germany)
      Account management, billing, provisioning

        Stripe (payment processing)
          EU entity for EU customers, US entity for NA customers

        Tenant VPS (customer's chosen region)
            EU: Netcup Germany/Austria
            NA: Netcup Manassas, Virginia (US)

            Tool containers (CRM, email, files, etc.)
               All data stored locally on VPS disk

            OpenClaw (AI runtime)
               Session transcripts stored locally

            Safety Wrapper Extension
```

```
  Secrets registry (encrypted SQLite)
  Audit log (append-only)
  Outbound redaction layer


[REDACTED prompts only]


OpenRouter → LLM providers
  (Anthropic, Google, DeepSeek, etc.)
  Prompts contain NO secrets, NO raw credentials
  Configurable PII scrubbing before transmission
```

1.3 3. GDPR Compliance

1.3.1 3.1 LetsBe's Role Under GDPR

For customer account data (Hub): LetsBe is the **data controller**. We determine the purposes and means of processing account data (name, email, billing).

For customer business data (tenant VPS): LetsBe is the **data processor**. The customer is the controller — they decide what data their CRM, email, files, and AI agents contain. We process it on their behalf to deliver the service.

For AI inference data: LetsBe is a processor, and LLM providers are **subprocessors**. The customer's data (redacted) passes through LetsBe's infrastructure to the LLM provider for inference.

1.3.2 3.2 Legal Basis for Processing

Processing Activity	Legal Basis (Art. 6)	Notes
Account creation and management	Art. 6(1)(b) — Contract performance	Necessary to deliver the service
Payment processing	Art. 6(1)(b) — Contract performance	Necessary for billing
Server provisioning and maintenance	Art. 6(1)(b) — Contract performance	Core service delivery
AI agent processing of customer data	Art. 6(1)(b) — Contract performance	The customer instructs the AI — we execute

Processing Activity	Legal Basis (Art. 6)	Notes
LLM inference (sending redacted prompts)	Art. 6(1)(b) — Contract performance	Essential for AI functionality
Token usage telemetry	Art. 6(1)(f) — Legitimate interest	Billing accuracy, abuse prevention, service optimization
Error and performance monitoring	Art. 6(1)(f) — Legitimate interest	Service reliability
Marketing emails (post-signup)	Art. 6(1)(a) — Consent	Opt-in only, unsubscribe available
Cookie analytics (website)	Art. 6(1)(a) — Consent	Cookie banner with granular consent

1.3.3 3.3 Data Subject Rights Implementation

GDPR grants data subjects (the customer and their end users) specific rights. Here's how the platform supports each one:

Right	Article	How We Implement It
Right of Access (Art. 15)	Customer can export all data from their VPS at any time via tool UIs (CRM export, file download, email export). Hub account data available via customer portal.	Response within 30 days.
Right to Rectification (Art. 16)	Customer has full admin access to all tools on their VPS. They can edit any data directly. Hub account data editable in customer portal.	Self-service — no ticket needed.

Right	Article	How We Implement It
Right to Erasure (Art. 17)	Customer can delete any data from their tools. Full account deletion: Hub removes account data, VPS is wiped and deprovisioned. See §3.6 for deletion procedures.	VPS wipe is irreversible — confirmed before execution.
Right to Restriction (Art. 18)	Customer can disable specific AI agents, restrict tool access, or set autonomy to Level 1 (read-only). Hub can freeze an account (stops all AI processing).	Granular per-agent and per-tool controls.
Right to Data Portability (Art. 20)	All tools on the VPS are open-source with standard export formats (CSV, JSON, MBOX, CalDAV, WebDAV). No vendor lock-in. AI conversation history exportable as JSON/Markdown.	Customer owns the server — they can SSH in directly.
Right to Object (Art. 21)	Customer can object to AI processing specific data categories. Safety Wrapper can be configured to exclude certain data types from AI context. Marketing emails have one-click unsubscribe.	Configurable per-agent data access rules.

Right	Article	How We Implement It
Automated Decision-Making (Art. 22)	AI agents propose actions — they do not make binding decisions without human oversight. Autonomy levels (§6 of Technical Architecture) ensure human approval for consequential actions.	No fully automated decisions affecting legal rights.

1.3.4 3.4 Data Processing Agreement (DPA)

LetsBe provides a standard DPA to all customers. The DPA covers Article 28 requirements:

DPA Element	Content
Subject matter and duration	Processing customer business data via AI-powered tool management, for the duration of the subscription
Nature and purpose	Storage, retrieval, AI-assisted analysis, and automated management of business data across customer-selected tools
Type of personal data	Contact records, email content, calendar events, file contents, invoicing data, website analytics — as determined by customer's tool selection
Categories of data subjects	Customer's employees, clients, contacts, website visitors — as determined by customer's use of tools
Controller obligations	Customer determines what data enters the platform, configures AI autonomy levels, manages user access
Processor obligations	LetsBe provides infrastructure, maintains security measures, processes data only on documented instructions, assists with data subject requests, notifies of breaches within 72 hours

DPA Element	Content
Subprocessors	Listed in §5 — customer has right to object to new subprocessors with 30 days notice
International transfers	Detailed in §4 — SCCs and adequacy decisions as applicable
Technical and organizational measures (TOMs)	Detailed in §6
Data return and deletion	Upon termination: customer has 30 days to export data, after which VPS is securely wiped
Audit rights	Customer may request evidence of compliance; LetsBe provides SOC 2 report (when available) or equivalent documentation

1.3.5 3.5 Records of Processing Activities (ROPA)

GDPR Article 30 requires maintaining records of processing activities. LetsBe maintains two ROPAs:

Controller ROPA (for Hub/account data): - Processing activity, purpose, legal basis, categories of data subjects, categories of data, recipients, international transfers, retention periods, TOMs reference

Processor ROPA (for tenant data processed on behalf of customers): - Categories of processing per customer, subprocessors involved, international transfers, TOMs reference

These records are maintained internally and available to supervisory authorities on request.

1.3.6 3.6 Data Retention & Deletion

Data Category	Retention Period	Deletion Method
Active tenant VPS data	Duration of subscription	Customer manages directly
Tenant VPS after cancellation	30 days grace period for data export	Secure VPS wipe (full disk overwrite via hosting provider API + VPS deletion)
Hub account data	Duration of subscription; soft-deleted at termination, hard-deleted after backup rotation (90 days)	Database soft-delete + backup rotation
Hub billing records	7 years (legal/tax obligation per German HGB §257)	Automated purge after retention period
AI conversation transcripts	Duration of subscription (on tenant VPS)	Deleted with VPS wipe

Data Category	Retention Period	Deletion Method
Token usage telemetry	24 months (aggregated, no PII)	Automated purge
Support tickets	24 months after resolution	Automated purge
Audit logs (tenant VPS)	Duration of subscription	Deleted with VPS wipe
Backups (Netcup snapshots)	7 daily snapshots, rolling	Oldest snapshot auto-deleted when new one is created

Deletion procedures:

1. **Customer requests account deletion** → Hub marks account for deletion → sends confirmation email → 48-hour cooling-off period
2. **After cooling-off** → Hub notifies customer that 30-day export window begins → customer can download all data via VPS tools and SSH access
3. **After 30-day window** → Hub triggers VPS deprovisioning via Netcup API → VPS disk is wiped → VPS instance deleted → all snapshots deleted
4. **Hub data** → Account record soft-deleted → billing records retained per legal obligation → all other data purged → soft-deleted record hard-deleted after backup rotation (90 days)

1.3.7 3.7 Breach Notification

Detection: The Safety Wrapper logs all tool executions, credential accesses, and anomalous patterns. The Hub monitors tenant health and connectivity. Unusual patterns (mass data export, credential access spikes, unauthorized API calls) trigger alerts.

Notification timeline: - **Internal:** Security team notified immediately upon detection - **Supervisory authority:** Within 72 hours of becoming aware (GDPR Art. 33) - **Affected customers:** Without undue delay if breach poses high risk to rights and freedoms (GDPR Art. 34) - **Affected data subjects:** As directed by the customer (controller) for breaches affecting their tool data

Breach response plan: 1. Contain — isolate affected VPS, revoke compromised credentials 2. Assess — determine scope, data categories affected, number of data subjects 3. Notify — supervisory authority (72h), customer (without undue delay), data subjects (if high risk) 4. Remediate — patch vulnerability, rotate all affected credentials, update security measures 5. Document — full incident report with timeline, impact assessment, remediation steps 6. Review — post-incident review within 14 days, update security procedures

---

1.4 4. International Data Transfers

1.4.1 4.1 Data Residency by Region

Customers choose their data center region at signup. Each region is served by Netcup infrastructure:

Region	Data Center Location	Jurisdiction	Default For
EU	Netcup — Nuremberg, Germany / Vienna, Austria	EU (GDPR applies natively)	European customers
NA	Netcup — Manassas, Virginia, USA	US (CCPA/state laws apply)	North American customers

EU region: Customer business data does not leave the EU. This eliminates the need for cross-border transfer mechanisms for the vast majority of data processing.

NA region: Customer business data stays in the US. North American customers benefit from lower latency (~20ms vs ~100ms+). CCPA/state privacy laws apply. NA customers who prefer GDPR-level protections can opt into the EU region instead.

Note: The Hub (account management, billing) always runs in the EU regardless of the customer's VPS region. Netcup pricing varies slightly by region (approximately ±€1-2/mo depending on server tier).

1.4.2 4.2 LLM Inference — The One Cross-Border Flow

The only data that regularly leaves the EU is **redacted AI prompts** sent to LLM providers for inference. This data:

- Has all secrets, credentials, and API keys stripped by the Safety Wrapper
- Has configurable PII scrubbing (can be enabled per customer or per agent)
- Is transient — LLM providers process it for inference and do not retain it for training (verified per provider policy and DPA)
- Contains business context (task descriptions, tool outputs) but not raw credentials

Transfer mechanisms by provider:

LLM Provider	Data Center	Transfer Mechanism	Training on Customer Data
Anthropic (Claude)	US	EU-US Data Privacy Framework + SCCs	No (per API terms)
Google (Gemini)	EU + US	EU-US Data Privacy Framework + SCCs	No (per API terms, when using paid API)
DeepSeek	China	SCCs + supplementary measures + enhanced redaction	No (per API terms) — extra scrutiny required
OpenRouter (aggregator)	US	EU-US Data Privacy Framework + SCCs	No (passthrough only, per DPA)

DeepSeek special handling: Given the geopolitical sensitivity of data transfers to China, LetsBe implements enhanced measures for DeepSeek routes: - Maximum redaction level enabled by default (PII scrubbing mandatory) - Customer opt-in required (not enabled by default) - Transparent disclosure in the model selection UI: "This model is hosted in China. Enhanced privacy protections are applied." - Customer can block specific providers entirely via settings

1.4.3 4.3 EU-US Data Privacy Framework

The EU-US Data Privacy Framework (DPF), adopted July 2023, provides an adequacy decision for transfers to certified US organizations. LetsBe verifies that US-based subprocessors (Stripe, Anthropic, OpenRouter) participate in the DPF. As a fallback, Standard Contractual Clauses (SCCs, 2021 version) are included in all subprocessor DPAs.

1.4.4 4.4 North American Customers

North American customers can choose between two regions:

- **NA region (Manassas, Virginia):** Lower latency for US/Canadian users. Data is subject to US jurisdiction. CCPA and applicable state privacy laws apply. LetsBe still applies its full security architecture (isolated VPS, secrets firewall, encryption at rest) regardless of region.
- **EU region (Germany/Austria):** Available as an opt-in for NA customers who prefer GDPR-level protections. Higher latency but stronger regulatory protections.

In either case, the Hub (account management) runs in the EU, so billing and account data are always GDPR-protected. The customer's VPS region is selected at provisioning and cannot be changed without re-provisioning (data migration assistance available).

1.5 5. Subprocessor Management

1.5.1 5.1 Current Subprocessors

Subprocessor	Purpose	Data Processed	Location	DPA Status
Netcup GmbH	VPS hosting	All tenant data (encrypted at rest on VPS)	Germany, Austria (EU region); Manassas, Virginia (NA region)	DPA available via Netcup CCP
OpenRouter	LLM API aggregation	Redacted AI prompts (transit only)	US	DPA required — verify DPF certification

Subprocessor	Purpose	Data Processed	Location	DPA Status
Anthropic	LLM inference (Claude models)	Redacted AI prompts (transit only)	US	API terms prohibit training; DPA available
Google	LLM inference (Gemini models)	Redacted AI prompts (transit only)	EU + US	API terms prohibit training (paid tier); DPA available
DeepSeek	LLM inference (DeepSeek models)	Redacted AI prompts (transit only, max redaction)	China	DPA + SCCs + supplementary measures required
Stripe	Payment processing	Customer name, email, payment method	EU (for EU customers), US (for NA customers)	DPA included in Stripe Terms
Poste Pro (self-hosted)	Transactional emails from Hub	Customer email address, email content	Self-hosted on LetsBe infrastructure (Hub server)	N/A — not a third-party subprocessor. If a relay service is adopted, add here with 30-day notice.

1.5.2 5.2 Subprocessor Change Process

Per GDPR Article 28(2), customers have the right to be informed of and object to new subprocessors:

1. **30-day advance notice** — LetsBe publishes new subprocessor additions on a changelog page and notifies customers via email
2. **Objection window** — Customers have 30 days to object on reasonable data protection grounds
3. **Resolution** — If objection cannot be resolved, customer may terminate without penalty within the objection window
4. **DPA flow-down** — All new subprocessors must sign DPAs with equivalent protections before processing begins

1.5.3 5.3 LLM Provider Vetting

Before adding a new LLM provider, LetsBe verifies:

- **No training on customer data** — confirmed via API terms, DPA, or written commitment

- **Data retention** — provider does not retain prompts or completions beyond the inference request (or has a clear, short retention window for abuse monitoring only)
- **Transfer mechanism** — valid adequacy decision, DPF certification, or SCCs in place
- **Security certifications** — SOC 2, ISO 27001, or equivalent
- **Breach notification** — provider commits to notifying LetsBe of breaches without undue delay

---

1.6 6. Technical & Organizational Measures (TOMs)

These measures implement GDPR Article 32 (security of processing) and form the basis for Annex II of the DPA.

1.6.1 6.1 Encryption

What	Method	Key Management
Data at rest (VPS disk)	Netcup full-disk encryption (provider-managed)	Netcup infrastructure
Secrets registry	AES-256-CBC with scrypt key derivation	Key generated at provisioning, stored on VPS filesystem (not in AI context)
Data in transit (user ↔ Hub)	TLS 1.3 (HTTPS)	Let’s Encrypt certificates, auto-renewed
Data in transit (user ↔ tenant VPS)	TLS 1.3 via nginx reverse proxy	Let’s Encrypt certificates, auto-renewed
Data in transit (Safety Wrapper ↔ LLM)	TLS 1.3 (HTTPS to OpenRouter)	OpenRouter TLS certificates
Backups (Netcup snapshots)	Provider-encrypted snapshots	Netcup infrastructure
SSH access	ED25519 keys, port 22022	Key-only auth, no password login

1.6.2 6.2 Access Control

Control	Implementation
Customer access to VPS tools	Keycloak SSO — single sign-on across all tools
Customer access to Hub	Email + password, session-based auth

Control	Implementation
Admin access to Hub	Role-based access control (Prisma + middleware)
SSH access to VPS	Key-only, port 22022, fail2ban (5 attempts → 300s ban)
AI agent access to tools	Per-agent tool allow/deny lists (OpenClaw config)
AI agent operational scope	Three-tier autonomy levels with command gating (Safety Wrapper)
Inter-tenant isolation	Separate VPS per customer — no shared infrastructure beyond Hub
Tool container isolation	Per-tool Docker networks with fixed subnets

1.6.3 6.3 Secrets Management

Measure	Description
Credential generation	50+ unique credentials generated per tenant at provisioning (env_setup.sh)
Credential storage	Encrypted SQLite registry on VPS — never transmitted to LLMs
Credential rotation	Registry supports rotation with audit trail
Outbound redaction	All LLM-bound traffic passes through 4-layer redaction (registry match → placeholder substitution → regex safety net → heuristic detection)
Transcript redaction	`tool_result_persist` and `before_message_write` hooks strip secrets from stored transcripts
Side-channel credential exchange	User-provided secrets never enter AI conversation — exchanged via direct Safety Wrapper API

1.6.4 6.4 Network Security

Measure	Description
Firewall	UFW — only ports 80, 443, 22022 open
OpenClaw binding	Localhost only — not accessible from outside VPS
Safety Wrapper binding	Localhost only — only OpenClaw and Hub (via nginx) can reach it
Tool container networking	Per-tool isolated Docker networks (172.20.X.0/28), exposed via 127.0.0.1:30XX
SSRF protection	Browser tool has configurable domain allowlists
Rate limiting	OpenClaw: 10 attempts/60s with 300s lockout; Hub API: rate-limited endpoints

Measure	Description
DDoS protection	Netcup infrastructure-level protection + nginx rate limiting

1.6.5 6.5 Monitoring & Audit

Measure	Description
Audit log	Append-only log of all AI agent actions on tenant VPS
Token metering	Per-agent, per-model token counts reported to Hub
Hub telemetry	Aggregated metrics (no PII) — uptime, error rates, usage patterns
Backup monitoring	AI-monitored backup-status.json with automated alerting
Uptime monitoring	Uptime Kuma on each VPS + Hub-level health checks

1.6.6 6.6 Physical Security

Delegated to hosting provider (Netcup): - ISO 27001 certified data centers in Germany, Austria, and Manassas, Virginia (US) - TÜV Rheinland annual security audits - Controlled physical access, CCTV, security personnel - Redundant power supply, climate control, fire suppression - Multiple redundant network connections

1.6.7 6.7 Organizational Measures

Measure	Description
Data protection awareness	Founder-led (Matt) for now; formal training program when team grows
Incident response plan	Defined in §3.7 — detection, containment, notification, remediation
Vendor assessment	All subprocessors vetted for GDPR compliance, DPAs in place
Privacy by design	Architecture decisions (isolated VPS, secrets redaction, local storage) baked into the platform from day one
Data minimization	Hub stores only what’s needed for account management; all business data stays on tenant VPS

1.7 7. EU AI Act Compliance

The EU AI Act entered into force August 1, 2024, with obligations phasing in through August 2027. LetsBe must assess its obligations under this framework.

1.7.1 7.1 AI System Classification

The AI Act classifies AI systems by risk level. LetsBe's AI agents need to be assessed:

AI Act Category	Applicability to LetsBe	Rationale
Prohibited (Art. 5)	Not applicable	LetsBe does not perform social scoring, real-time biometric identification, emotional inference in workplaces, or other prohibited practices
High-risk (Annex III)	Likely not applicable for V1	LetsBe agents manage business tools — they do not make employment decisions, assess creditworthiness, or perform other Annex III high-risk functions. If customers use CRM/sales tools for automated lead scoring that affects access to services, this needs monitoring.
Limited-risk (transparency)	**Applicable**	AI agents interact with users — transparency obligations apply
Minimal-risk	Most functionality falls here	General business automation, scheduling, file management

1.7.2 7.2 Transparency Obligations (Art. 50)

As a provider of an AI system that interacts with humans, LetsBe must:

Obligation	Implementation
Inform users they're interacting with AI	The product is explicitly marketed as AI agents. Every agent is labeled as AI in the UI. The onboarding flow introduces agents as "your AI team."
AI-generated content disclosure	When AI agents send external communications (email, chat), the External Communications Gate (Decision #30) requires human review and approval. Outbound messages include a configurable disclosure footer.
Synthetic content marking	Not applicable for V1 — agents don't generate deepfakes or synthetic media

1.7.3 7.3 GPAI Model Obligations

LetsBe is a **deployer** of general-purpose AI models (Claude, Gemini, DeepSeek), not a **provider**. The provider obligations (technical documentation, training data transparency, systemic risk assessment) fall on Anthropic, Google, DeepSeek respectively.

As a deployer, LetsBe's obligations are: - Use GPAI models in accordance with their intended purpose and instructions for use - Maintain transparency about which models are being used (disclosed in advanced settings) - Implement human oversight measures (autonomy levels, command gating) - Monitor for incidents and report to providers and authorities as required

1.7.4 7.4 AI Literacy (Art. 4)

Effective February 2, 2025, all organizations deploying AI must ensure sufficient AI literacy among staff and users. LetsBe addresses this through:

- Clear, non-technical onboarding that explains what the AI can and cannot do
- Autonomy levels that let users control AI scope based on their comfort
- In-app explanations of AI actions ("I'm about to do X because Y — approve?")
- Documentation and help resources explaining AI capabilities and limitations
- The Dispatcher agent defaults to asking when intent is ambiguous rather than assuming

1.7.5 7.5 Record-Keeping

LetsBe maintains records relevant to AI Act compliance: - Audit logs of all AI agent actions (per-tenant VPS) - Token usage and model selection logs - Customer autonomy level configurations - AI incident reports (if any)

1.8 8. North American Privacy Compliance

1.8.1 8.1 CCPA/CPRA (California)

The California Consumer Privacy Act, as amended by the California Privacy Rights Act, applies to businesses meeting revenue or data processing thresholds. While LetsBe may not initially meet the $26.6M revenue threshold, building for CCPA compliance from day one is the right approach.

CCPA Right	LetsBe Implementation
Right to Know	Customer portal shows all collected data; VPS tools provide direct data access
Right to Delete	Account deletion flow (§3.6); tool-level data deletion self-service
Right to Opt-Out of Sale	LetsBe does not sell personal information. Period. No data brokers, no ad targeting, no third-party data sharing for marketing.
Right to Non-Discrimination	No service differences based on privacy choices
Right to Correct	Self-service editing in customer portal and VPS tools
Right to Limit Use of Sensitive PI	Configurable AI data access rules per agent

"Do Not Sell or Share" compliance: LetsBe's architecture inherently satisfies this — customer business data stays on their VPS and is never shared with third parties. Redacted LLM prompts are not "sold" or "shared" under CCPA definitions (they're processed for service delivery under the customer's instructions).

Automated Decision-Making (ADMT): Per CCPA's 2026 ADMT regulations, LetsBe's AI agents do not make decisions that "replace or substantially replace human decision-making" in ways that affect access to services, employment, or other significant categories. The autonomy level system ensures human oversight for consequential actions.

1.8.2 8.2 US State Privacy Law Patchwork

Multiple US states have enacted comprehensive privacy laws with varying requirements. LetsBe's approach: build to the strictest standard (currently CCPA/CPRA with 2026 ADMT rules), which covers the requirements of other state laws.

State Law	Effective	Key Difference from CCPA	LetsBe Compliance
Virginia CDPA	Jan 2023	No private right of action; 30-day cure	Covered by CCPA compliance
Colorado CPA	Jul 2023	Universal opt-out mechanism required	"Do Not Sell" not applicable (we don't sell data)

State Law	Effective	Key Difference from CCPA	LetsBe Compliance
Connecticut CTDPA	Jul 2023	Broader “sale” definition	N/A — no data sales
Indiana ICDPA	Jan 2026	Mirrors Virginia	Covered
Kentucky KCDPA	Jan 2026	Mirrors Virginia	Covered
Rhode Island RIDPA	Jan 2026	60-day cure period	Covered

1.8.3 8.3 Canadian PIPEDA

For Canadian customers, PIPEDA (Personal Information Protection and Electronic Documents Act) applies. LetsBe’s GDPR-compliant practices exceed PIPEDA requirements in most areas. Key considerations:

- Consent for collection (covered by our signup and DPA flow)
- Purpose limitation (data used only for service delivery)
- Data residency (NA region in Virginia for low latency; EU region available if preferred — adequacy decision between EU and Canada exists)
- Breach notification (72-hour timeline aligns with PIPEDA requirements)

1.9 9. AI-Specific Privacy Controls

1.9.1 9.1 Secrets Firewall

The most significant privacy control in the platform. Detailed in Technical Architecture §3.2.1 and §13. Key properties:

- Four-layer outbound redaction (registry match, placeholder substitution, regex safety net, heuristic detection)
- All 50+ provisioned credentials registered and tracked
- Pattern matching catches credentials the registry might miss
- AI never sees raw credential values — only deterministic placeholders like `[REDACTED:postgres_pa`
- Side-channel credential exchange for user-provided secrets
- Non-bypassable — runs at the transport layer, not dependent on AI behavior

1.9.2 9.2 Configurable PII Scrubbing

Beyond credential redaction, the Safety Wrapper supports configurable PII scrubbing before LLM inference:

PII Category	Default	Configurable
Credentials and API keys	Always scrubbed	Cannot be disabled
Email addresses in tool outputs	Off (needed for most tasks)	Customer can enable
Phone numbers in tool outputs	Off	Customer can enable
Physical addresses	Off	Customer can enable
Financial data (invoice amounts, etc.)	Off	Customer can enable
Names in tool outputs	Off	Customer can enable

Trade-off: More scrubbing = more privacy, but less useful AI. A marketing agent that can't see email addresses can't draft personalized emails. Defaults are set for maximum utility with mandatory credential protection. Customers in regulated industries (healthcare, legal) can dial up scrubbing.

1.9.3 9.3 AI Conversation Data Handling

Property	Implementation
Storage location	Tenant VPS only (JSON files on local disk)
Encryption	Protected by VPS disk encryption
Retention	Duration of subscription — auto-pruned per OpenClaw defaults (30-day stale session cleanup)
AI access to history	Per-agent memory search with configurable scope
Export	JSON/Markdown export via customer portal or direct SSH
Deletion	Customer can delete individual conversations or all history
Transcript redaction	`before_message_write` hook strips secrets before session persistence

1.9.4 9.4 External Communications Gate (Decision #30)

When AI agents send external communications (emails, chat messages), an independent safety layer applies:

- All outbound external communications require human approval regardless of autonomy level
- Each message shows: recipient, subject, full content preview
- Customer can approve, edit, or reject
- This prevents the AI from inadvertently sharing sensitive data in external communications

- Configurable: customers can whitelist specific communication types after building trust

---

1.10 10. Security Certifications Roadmap

1.10.1 10.1 Current State (Pre-Launch)

- Netcup: ISO 27001 certified data centers, TÜV Rheinland audited
- Stripe: PCI DSS Level 1, SOC 2 Type 2
- Anthropic: SOC 2 Type 2
- LetsBe itself: No certifications yet — pre-revenue startup

1.10.2 10.2 Planned Certifications

Certification	Target Timeline	Why
SOC 2 Type 1	Year 1 post-launch	Baseline security certification — expected by B2B customers
SOC 2 Type 2	Year 1-2 post-launch	Demonstrates sustained security practices over audit period
ISO 27001	Year 2-3	International standard — important for EU enterprise customers
GDPR certification (Art. 42)	When available	Voluntary certification mechanism under GDPR — still emerging

1.10.3 10.3 Interim Measures

Until formal certifications are obtained: - Published security page with TOMs, architecture overview, and FAQ - DPA available to all customers - Subprocessor list maintained

and updated - Security questionnaire responses (CAIQ framework) available on request - Penetration testing (planned before launch, annual thereafter) - Vulnerability disclosure program

---

1.11 11. Customer-Facing Security Artifacts

1.11.1 11.1 Published Security Page

A public page on the LetsBe website covering: - Architecture overview (isolated VPS, secrets firewall, four-layer security) - Data residency (EU and NA data center options) - Encryption standards - Subprocessor list with update history - Compliance status (GDPR, AI Act, CCPA) - Contact for security questions

1.11.2 11.2 DPA (Available on Request, Self-Service Preferred)

Pre-signed DPA available in the customer portal. Customers accept it as part of signup (checkbox). Enterprise customers can request custom DPA negotiations.

1.11.3 11.3 Security FAQ for Sales

Common questions and answers for sales conversations:

"Where is my data stored?" On your own dedicated server in the region you choose — either Nuremberg, Germany (EU) or Manassas, Virginia (US). European customers default to the EU region; North American customers default to the NA region for lower latency. Your business data stays in your chosen region. Only AI prompts (with all secrets removed) are sent to AI providers for processing.

"Can you access my data?" We have SSH access to your server for maintenance and support. This access is logged and auditable. We never access your data for purposes other than service delivery and support. You can revoke our SSH access if you prefer fully self-managed operation (advanced users only).

"Does the AI train on my data?" No. We use API access to AI providers (Anthropic, Google, etc.) under terms that explicitly prohibit training on customer data. Your business data never enters any AI training pipeline.

"What happens if I cancel?" You get 30 days to export all your data (using the tools directly, or via SSH). After 30 days, your server is securely wiped and deleted. Billing records are retained for 7 years per German tax law (HGB §257), since LetsBe operates from the EU.

"Are you GDPR compliant?" Yes. Our architecture is privacy-by-design: isolated servers in your chosen region (EU or NA), secrets that never leave your server, and a full DPA covering our processing activities. EU-region customers get native GDPR jurisdiction. NA-region customers can opt into the EU region if they prefer GDPR protections. We maintain records of processing, support all data subject rights, and have a documented breach notification process.

“What about the EU AI Act?” We classify as a deployer of general-purpose AI models, not a provider. Our transparency obligations are met through clear AI labeling, human oversight via autonomy levels, and external communications gating. We monitor regulatory developments and will adapt as requirements evolve.

“Do you have SOC 2?” Not yet — we’re a pre-launch startup. Our hosting provider (Netcup) has ISO 27001 certified data centers in both EU and US regions, and our AI provider (Anthropic) has SOC 2 Type 2. We plan to obtain SOC 2 Type 1 within our first year post-launch.

1.12 12. Implementation Priorities

1.12.1 12.1 Must-Have Before Launch

- ☐ DPA template finalized and available in customer portal
- ☐ Privacy Policy published (website + app)
- ☐ Terms of Service with data processing clauses
- ☐ Cookie consent banner on website (granular consent)
- ☐ Subprocessor list published
- ☐ Security page published
- ☐ Breach notification procedure documented and tested
- ☐ Data deletion procedure documented and tested
- ☐ Secrets firewall operational and tested
- ☐ PII scrubbing configurable per customer
- ☐ External Communications Gate operational
- ☐ Audit logging active on all tenant VPS instances
- ☐ Records of Processing Activities (ROPA) created

1.12.2 12.2 Within 6 Months Post-Launch

- ☐ Penetration test completed by third-party firm
- ☐ Data Protection Impact Assessment (DPIA) for AI processing completed
- ☐ Security questionnaire (CAIQ) responses prepared
- ☐ Vulnerability disclosure program launched
- ☐ SOC 2 Type 1 audit initiated
- ☐ CCPA-specific disclosures added for California users (if threshold met)
- ☐ AI Act conformity self-assessment documented

1.12.3 12.3 Within 12 Months Post-Launch

- ☐ SOC 2 Type 1 obtained
- ☐ SOC 2 Type 2 audit cycle begun
- ☐ Annual penetration test
- ☐ DPA review and update based on customer feedback
- ☐ Subprocessor audit (verify all DPAs current)

☐ AI Act compliance review (ahead of August 2026 high-risk deadline)
☐ Privacy training program for new team members

1.13 13. Open Questions

#	Question	Status	Notes
1	Data Protection Officer (DPO) appointment	Open	Required under Art. 37 if processing “on a large scale.” Assess once customer base reaches ~100 tenants. Matt may serve as interim DPO.
2	DPIA for AI-assisted business management	Open	Likely required for AI agents processing personal data across multiple tools. Complete before launch or within 6 months.
3	Supervisory authority registration	Open	Determine lead supervisory authority based on LetsBe’s EU establishment. Likely BfDI (Germany) given Netcup hosting.
4	EU representative appointment (Art. 27)	Open	Required if LetsBe is not established in the EU but offers services to EU residents. Depends on corporate structure.
5	Transactional email provider selection	Open	Choose EU-based provider to avoid cross-border transfer complexity.

#	Question	Status	Notes
6	DeepSeek transfer mechanism	Open	SCCs + supplementary measures need legal review given China data transfer complexity. May defer DeepSeek support until proper legal framework is in place.
7	Cookie analytics tool	Open	Select privacy-friendly analytics (likely Umami, already in tool stack — self-hosted).
8	Cyber insurance	Open	Evaluate coverage for data breach liability. Recommended before taking paying customers.

1.14 14. Changelog

Version	Date	Changes
1.0	2026-02-26	Initial framework. Data classification, GDPR compliance, international transfers, subprocessor management, TOMs, EU AI Act assessment, North American compliance, AI-specific privacy controls, security certification roadmap, customer-facing artifacts, implementation priorities.

Version	Date	Changes
1.1	2026-02-26	Added dual-region data center support (EU: Nuremberg/Vienna, NA: Manassas, Virginia). Updated data residency tables, data flow diagram, subprocessor entries, physical security references, Section 4.1/4.4, PIPEDA section, security page scope, and sales FAQ to reflect customer region choice.

This document is a living framework. It will be updated as regulations evolve, the platform matures, and customer requirements emerge. Legal counsel should review before finalizing the DPA, Privacy Policy, and Terms of Service.