Brainverse runs on certified infrastructure and allocates accountability between four parties. **Brainverse** operates isolation architecture, hook enforcement, and agent-definition discipline. **The client** owns the repository, the identity surface, and the data classification decisions. **Anthropic** operates the LLM inference layer under its published data policy. **Substrate providers** (GitHub, Railway) operate the certified storage, hosting, and database layers we build on. Each party is accountable for the controls below.

Control Category	Brainverse AGENT-TEAM OPERATOR	Client DATA OWNER	Anthropic LLM PROVIDER	Substrate Provider GITHUB / RAILWAY
Data Protection AT REST & IN TRANSIT	Enforces isolation architecture across repo RBAC, filesystem, and BrainSync memory scopes. Operates per-tenant boundaries with no shared runtime.	Classifies data (confidential by default), selects deployment mode, and controls what content is authorized for agent access.	Processes inference under Anthropic's published data policy. Zero-data-retention available directly or via Bedrock / Azure OpenAI.	Encrypts at rest (AES-256) and in transit (TLS 1.2+). Each provider publishes certifications at its own trust center.
Identity & Access RBAC, MFA, SSO	Enforces least privilege on staff accounts with hardware-backed MFA. No standing access to client data; support access is per-incident, time-boxed, logged.	Owns client admin accounts on GitHub, Anthropic, and the BrainSync dashboard. Client enables MFA / SAML on its own tenants.	Operates Anthropic console authentication, API key lifecycle, and enterprise SSO per client agreement.	Provides RBAC primitives (GitHub teams, Railway project roles). Client configures them.
Secrets & Keys STORAGE, ROTATION	Stores secrets in Railway Variables. Never in code or memory files. Rotates on personnel changes.	Delivers credentials over encrypted channels and revokes any Brainverse-held credential from the provider console at any time.	Issues and rotates API keys on client request. Enforces key scoping and revocation.	Operates the underlying secret store (Railway Variables and GitHub encrypted secrets) with its own KMS and audit trail.
Availability & Uptime PLATFORM SLAS	Publishes target RTO / RPO for BrainSync (targets, not contractual SLAs). Client-hosted and local-only modes operate without Brainverse in the loop.	Selects deployment mode matching its uptime needs. For client-hosted, controls its own availability envelope.	Operates Claude API availability under Anthropic's published SLA.	Each substrate publishes its own availability SLA (GitHub, Railway). SOC 2 Type II audited.
Patching & Vulnerabilities CVE RESPONSE	Reviews substrate advisories weekly. Target windows: 72 hours CRITICAL, 14 days HIGH, 30 days MEDIUM. Pins MCP server versions.	Patches its own GitHub repo dependencies where the client owns the delivery repo, and any client-hosted infrastructure.	Patches Claude API and model-serving infrastructure. Publishes model and policy changes on release.	Patches managed substrate infrastructure. Notifies on platform advisories.
Monitoring & Logging AUDIT TRAIL	Logs agent dispatches, tool invocations, memory writes, hook executions, and commit history in BrainSync. Hook-health telemetry monitored.	Owns access to logs for its own tenant. Opts in to enriched prompt / output logging where needed for its own compliance.	Provides Anthropic console logs (API calls, usage, errors) under the client's workspace or the operator's workspace per contract.	Provides provider-native audit logs (GitHub audit log, Railway logs). Retention varies per provider.
Incident Response NOTIFICATION, FORENSICS	Targets 72-hour notification for confirmed incidents affecting client data (target, not contractual). Rotates credentials, preserves evidence, issues written report.	Declares incidents on its own side, owns internal comms, sets its own regulatory notification cadence (GDPR, state AG, customer disclosure).	Notifies per Anthropic's published incident policy. Honors forensics requests under enterprise agreements.	Notifies per each provider's SOC 2 incident policy. Provides forensic data under support tickets.
Backup & Recovery BC / DR	Operates BrainSync DB backups on Railway's scheduled cadence. Git-based configuration is distributed by design.	Owns backup policy for client-hosted and local-only modes. Retains its own copy of the delivered GitHub repository.	N/A — stateless inference. No client data stored beyond configurable retention window.	Operates managed backups (GitHub, Railway) per each provider's published retention policy.
Compliance & Attestation SOC 2, ISO, HIPAA	Not certified. Publishes operational discipline, this matrix, and written responses to SIG / CAIQ. HIPAA BAA available case-by-case as custom Enterprise addendum.	Owns its own regulatory obligations (GDPR data controller role, HIPAA covered-entity status, SOX scoping, state privacy laws).	SOC 2 Type II, ISO 27001. Enterprise BAA available directly from Anthropic for HIPAA workflows.	GitHub Enterprise (SOC 1/2 Type II, ISO 27001, FedRAMP). Railway (SOC 2 Type II). Each publishes a current trust center.
Subprocessor Disclosure CHANGE NOTICE	Publishes full subprocessor list at brainverse.ai/security/subprocessors (HTML, CSV, JSON). Targets 30 days' notice before adding a subprocessor that processes client data.	Subscribes to change notifications and reviews new subprocessors against its own vendor management policy.	Publishes Anthropic's subprocessor list at anthropic.com. Governs its own upstream model and infrastructure providers.	Each provider publishes its own subprocessor list at its trust center. Brainverse tracks material substrate changes.