Brainverse for Your Company: What Your IT Team Should Know

Someone on your team is asking to deploy a Brainverse agent team for a workflow they spend too much time on. This document is what your IT and security lead needs to decide whether that is a reasonable request. It is not a sales pitch. It describes what Brainverse is, what data it touches, who can see that data, and what your team can control. If any of this raises a specific question your team needs answered, the contact inbox for security questions is at the bottom of page 3.

Brainverse deploys custom AI agent teams for operational workflows. The agent team runs inside a GitHub repository that your organization can own, under identity credentials your organization issues, against an inference endpoint (Anthropic's Claude API) that either sits on your team's account or on Brainverse's, depending on how you deploy. The employee who is asking for this is not asking you to install a vendor platform. They are asking you to approve a small, inspectable codebase running on infrastructure you already evaluate.

What data does this touch?

The agent team reads and writes only the data you explicitly give it access to. In practice that means: the repositories or folders you connect it to, the documents you upload or paste in, and the workflow notes the employee writes while working with the agent. It does not crawl your wider network, your other SaaS tools, or your internal systems unless you deliberately connect them through a named integration.

Prompts and outputs are sent to Anthropic's Claude API for inference. Anthropic's commercial API terms prohibit using customer inputs or outputs to train Anthropic's foundation models. If your team operates in Brainverse-hosted mode, the Brainverse account runs with Zero Data Retention enabled, which eliminates Anthropic's standard 30-day safety retention window. If your team operates in client-hosted or local-only mode against your own Anthropic workspace, your team configures Zero Data Retention on your own account, and Brainverse verifies the setting before the deployment goes live. Brainverse itself does not train, fine-tune, or aggregate on your data for any purpose.

Who can see it?

Access follows the deployment mode your team chooses. There are four:

> **Local-only mode.** The agent team runs on a laptop your team owns. No Brainverse-operated infrastructure is involved. The only outbound traffic is Anthropic inference. Brainverse staff have no access to data in this mode unless a specific support request is approved and scoped.

> **Client-hosted mode.** The agent team runs inside your AWS, Azure, or GCP environment. Your cloud, your identity access management, your network boundary. Brainverse staff operate inside the boundary your team defines, typically through a named access path that your team provisions and revokes on your schedule. Data does not traverse Brainverse-owned infrastructure.

Brainverse-hosted mode. The agent team runs on Brainverse-operated infrastructure (Railway, with client code in the client's GitHub organization). This is the fastest mode to deploy, and it means a shared-responsibility slice for data at rest sits with Brainverse. This mode is not compatible with HIPAA-regulated workflows; workflows that process Protected Health Information should select local-only or client-hosted mode instead.

Hybrid mode. Runtime in your cloud, dashboard and memory storage on Brainverse-hosted infrastructure. Useful when your team wants client-cloud runtime isolation without self-hosting the Brainverse dashboard.

Brainverse staff do not have standing administrative access to your production data in any mode. When a support request does require access, it is provisioned per incident, time-boxed, and logged on the underlying provider's own audit trail (GitHub, Railway, Anthropic, your cloud). Brainverse does not currently operate a unified cross-provider access log.

What your team controls

Your team holds every credential Brainverse touches. You can:

- Revoke the GitHub personal access token Brainverse holds for your engagement, at any time, from your own GitHub organization admin console.
- Revoke the Anthropic API key for your workspace, at any time, from the Anthropic console.
- Revoke the Brainverse dashboard authentication token, if you use the dashboard, from the dashboard admin interface.
- Revoke any OAuth grants Brainverse holds for integrated services (Google Workspace, Microsoft 365, GoHighLevel, other connected systems), at any time, from those providers' admin consoles.

Each revocation path is independent. Any one of them stops the corresponding capability within seconds. Your team does not need Brainverse's cooperation to revoke access. Offboarding a Brainverse engagement from your team's side is a checklist of provider-side revocations, each taking under 30 seconds on the relevant admin console.

Your team also controls scope. Brainverse's default runtime assumes the permission model Claude Code ships with, which sandboxes agent filesystem access to the deployed repository. Expanded access (additional filesystem paths, new MCP integrations, broader bash capability) is an explicit engagement decision and is documented in your engagement's configuration file.

Certifications we inherit

Brainverse runs on certified infrastructure. Anthropic maintains SOC 2 Type II and ISO 27001. GitHub Enterprise maintains SOC 1/2 Type II, ISO 27001, and FedRAMP. Railway maintains SOC 2 Type II. Each provider publishes their current certifications at their own trust center; links are in our shared responsibility matrix. Brainverse itself does not hold SOC 2, ISO 27001, or any formal certification. Our posture is operational discipline layered on top of certified infrastructure, with clear shared-responsibility allocation between Brainverse, your team, and the providers we run on.

If your procurement process requires a certified vendor for the party holding the data, the data itself lives on certified infrastructure your team can control (your GitHub organization, Anthropic under a direct or Bedrock relationship). If your procurement process requires a certified vendor for Brainverse itself, Brainverse is not the right fit today.

Honest limits

The following are admissions your team should know before approving a Brainverse deployment. Brainverse wants your IT lead to cite these in our voice rather than discover them later.

No Brainverse-level SOC 2 or ISO 27001 certification. Brainverse does not hold a formal security certification and does not operate a certification roadmap. Compensating context: the data Brainverse touches lives on certified providers your team can control directly.

No formal background-check program for staff. Brainverse is a small team. Staff screening is a reference-and-identity process documented in the engagement's security evidence folder, not a vendor-delivered background check. Compensating context: the same small-team posture means access grants are reviewed by a named individual, not approved through a ticketing queue that can be gamed.

No 24x7 Security Operations Center. Brainverse does not operate a staffed SOC. Compensating context: Brainverse targets notification within 72 hours of confirmed impact for security incidents affecting client data, consistent with GDPR Article 33's processor-to-controller expectation. Specific contractual SLAs are available at the Enterprise tier and are set in the Order Form.

No unified cross-provider access log for Brainverse staff actions. Access is logged in each underlying provider's audit trail (GitHub, Railway, Anthropic). Compensating context: your team can pull a full access log for any client-hosted resource from the provider directly, without Brainverse's cooperation.

No cryptographically signed, tamper-evident audit events on Brainverse's own session store. Session metadata is stored in Postgres over TLS. Compensating context: for evidence workflows that require tamper-evidence, the provider-side audit log (GitHub, Railway) is a stronger source than Brainverse's session table.

Prompt injection is not a solved problem in the field. Brainverse's defense is layered: a chain of fail-closed tool-call hooks blocks credential writes, destructive bash, and configuration tampering at the runtime layer. Data-interpretation within an agent's authorized scope is a mitigation, not a hard boundary. Compensating context: your team should treat agent output as advisory on sensitive decisions until the underlying source is verified.

What to do next

If your team's question is about architecture, a full deployment-modes matrix (egress allowlists, firewall compatibility, HIPAA posture per mode, training-data handling per mode) is at brainverse.ai/security/deployment-modes. The full shared responsibility matrix is a one-page PDF downloadable from brainverse.ai/security.

If your team's question is about a specific control, a full threat model and hook architecture document is at brainverse.ai/security/threat-model. An executive summary of the threat model is available on request via the form at brainverse.ai/security/threat-model-request.

If your team's question is one this document does not answer, send it directly to the inbox below. Brainverse targets acknowledgment within 5 business days for general inquiries and within 72 hours for confirmed incident notifications.

Security questions, vulnerability reports, questionnaire responses: **security@brainverse.ai**

Privacy requests, data subject rights, DPA questions: **privacy@brainverse.ai**

This document is maintained at brainverse.ai/security and may be re-downloaded at any time. Last updated 2026-04-17.