Superintelligent

SECURITY & PRIVACY PACK

AGENT TRANSFORMATION INTENSIVE

Information Security FAQ

Written for the security and IT teams reviewing this program: what it is, what it touches, what is shared, and how that data is protected.

QR Quick Reference

The essentials for a first-pass review. Each item is expanded in the sections that follow.

ITEM	SUMMARY
Program type	Six-week hands-on training program for your core AI team. Not software installed or run in your environment.
Superintelligent system access	None. Your employees build and run agents in your environment, using your approved tools.
Credentials shared with SI	None. SI never receives credentials, OAuth tokens, or API keys.
Data shared with SI	Participant account and progress data, plus any documents a participant chooses to upload.
Data hosting	Superintelligent infrastructure. Encrypted in transit (TLS 1.2+) and at rest (AES-256).
Data retention	60 days after program completion, then deleted. Deletion also available on request.
Certifications	ISO 27001 and NIST CSF aligned; SOC 2 Type 2 on roadmap, not yet certified.
Privacy role	SI acts as a data processor under GDPR and CCPA, processing client data only as defined in the client agreement.

0 Start Here

The Agent Transformation Intensive is a training program, not software that Superintelligent installs or operates inside your environment. Your team builds agents on its own work, using your company-approved coding tool, connecting to your own systems through connections that your IT and security teams provision and control.

> **HOW THIS DIFFERS FROM A STANDARD VENDOR REVIEW**
>
> A typical software vendor receives access to your systems or a copy of your data, and the review centers on how that vendor protects it. Here, the access to your systems stays with your own employees and your own approved tooling. Superintelligent is not in the path between an agent and the system it connects to, and never receives the credentials for those connections.

The one place data reaches us is the platform account data and any documents participants choose to upload, which can include sensitive organizational detail. Section 3 sets out exactly what that surface is, and Section 4 sets out how it is protected.

How to read this document

Section 1	What this program is, and is not
Section 2	What systems it touches, and who gets access
Section 3	What is shared with Superintelligent
Sections 4 to 7	Controls, compliance, posture, and how to begin

1 What this program is, and is not

Q What is the Agent Transformation Intensive?

A six-week, hands-on program that trains a core team of three to eight of your employees to build, govern, and scale a workforce of AI agents. Each week, participants learn a capability, build it on their own real work inside your environment, and capture it in an enterprise strategy document. The agents are the medium for learning. The capability, the governance, and the operating model are what the team keeps.

Q Is this software we install in our environment, or a training program?

A training program. There is no Superintelligent software agent, daemon, or service deployed into your systems. Participants build in your company-approved coding tool, which you have already selected and installed. Superintelligent provides curriculum, structured exercises, office hours, and a web platform where participants track progress and, optionally, store their work.

Q Does Superintelligent host, run, or operate any agents inside our systems?

No. The agents are built and run by your team, in your environment, on your infrastructure. Superintelligent does not operate them and has no runtime presence in your environment.

Q Does the program deploy autonomous agents into production?

No. The program produces working agents built on participants' real work, along with vetted designs for your prioritized use cases. Moving any of those to production is downstream engineering work that your team owns and performs on its own timeline. Graduation does not place agents into your production systems.

Who is responsible for what?

RESPONSIBILITY	OWNER
Selecting, approving, and procuring the coding tool	Your organization
Provisioning and approving system connections (OAuth, admin consent, MCP servers)	Your IT and security teams
Building, running, and governing the agents	Your participants
Curriculum, exercises, office hours, and the learning platform	Superintelligent
Protecting data held on the Superintelligent platform	Superintelligent

Q Where does the agent-building actually happen?

Inside your environment, in your approved coding tool, on the participant's machine or your provisioned infrastructure. The strategy documents are authored by your team about your organization. The only material that leaves your environment is what a participant chooses to upload to the platform.

2 What systems it touches, and who gets access

Q What systems will the agents connect to during the program?

In Week 3, participants connect an operations agent to common workplace systems: mail, chat, and calendar. For Google organizations this is Gmail, Google Calendar, and Slack. For Microsoft organizations this is Outlook mail and calendar and Teams, through Microsoft Graph. Later weeks may connect read-only to a system of record, such as a CRM or PLM, only if a participant's use case calls for it and your team approves it.

Q Who provisions and owns those connections?

Your IT and workspace administrators. Each connection is made through your own OAuth flow or admin consent, scoped to the permissions your administrators grant, and follows least privilege. Your team decides what each agent can reach. Superintelligent does not create, broker, or hold these connections.

Q Does Superintelligent ever receive our credentials, tokens, or API keys?

No. Credentials, OAuth tokens, and connection configurations reside in the participant's approved coding tool and on the participant's machine or your infrastructure. They are not transmitted to Superintelligent and are not stored on the platform.

Q What coding tool is used, and is it enterprise-approved?

Each participant enters the program with an active, company-approved agentic coding tool already installed. Selecting and approving that tool is your decision and a precondition for joining, not part of the program. The program supports the major enterprise tools, including Cursor, Claude Code, Codex, GitHub Copilot, and Antigravity.

The corporate version of the program defaults to your approved tool. A separate open-source tool path exists in the platform but is gated off unless an administrator deliberately enables it, which keeps unapproved tooling out of the default experience.

Q Where does the model inference run, and who is the model provider?

Inference runs through your chosen coding tool and its model provider, under your own enterprise agreement with that provider. Superintelligent is not in that path and does not process or receive the model traffic generated as participants build.

Q How are third-party connectors and MCP servers handled?

Any connector or MCP server used to reach a system is subject to your security team's review and allowlisting before it is used. The program itself teaches participants how to run that risk review, so the practice is applied during the program rather than bypassed.

Q Will agents touch our systems of record, such as a CRM or PLM?

Only if you choose to allow it. Where a use case calls for it, the recommended pattern is read-only access through a connector your IT team registers, scoped to least privilege. This is your decision to grant, system by system, and is not required to complete the core program.

3 What is shared with Superintelligent

Q What does the Superintelligent platform collect and store?

Two categories of data:

- **Account and program data:** participant name and email, the organization and cohort, the selected tool path, progress and completion records, and notifications.
- **Documents a participant chooses to upload:** files attached to weekly tasks, and the strategy artifacts and context documents produced during the program.

Q What is in the strategy artifacts, and how sensitive is it?

These are the participant's own work product, written about your organization. Depending on how a participant writes them, they can contain sensitive detail, including named people, descriptions of your security posture, and tool-access levels. We treat this as the most sensitive category of data in the program, which is why the controls in Section 4 and the upload accommodation below apply specifically to it.

Q Are we required to upload sensitive documents to the platform?

No. Completing the program does not require uploading sensitive organizational documents. A participant can progress while keeping the source material in your own environment, so the platform records that the work was completed without holding the document itself. This is an operational accommodation agreed at the program level rather than an enforced platform mode, so your team should treat it as a documented agreement, not a technical control.

WHAT IS NOT SHARED WITH SUPERINTELLIGENT

- ✕ Access to your systems
- ✕ Credentials, tokens, or API keys
- ✕ Connected mailbox, chat, or calendar contents
- ✕ Production data from your systems of record
- ✕ The agents themselves
- ✕ Model inference traffic

Q Is our data used to train AI models?

No. Superintelligent does not use client data to train public AI models.

Q Does Superintelligent sell or share our data?

No. Superintelligent processes client data solely for the purposes defined in the client agreement and does not sell or share it.

Q Is our cohort space isolated from other companies?

Yes. The corporate version runs as a single-tenant cohort space for your organization. There is no cross-company community and no third-party messaging channel such as WhatsApp in the corporate version. Within your cohort, uploaded documents are visible only to the participant and their team members.

4 How Superintelligent protects the data it holds

The controls below apply to the data described in Section 3: account and program data, and any documents a participant uploads. All program data is held in Superintelligent's own infrastructure.

CONTROL AREA	WHAT WE DO
Security governance	An Information Security Management System aligned to ISO 27001 and the NIST Cybersecurity Framework. CEO executive oversight; the CTO maintains and reviews the program and coordinates assessments.
Access control	Least privilege throughout: role-based access management, multi-factor authentication for administrative accounts, regular access reviews, and removal of unused credentials. Staff sign confidentiality agreements.
Encryption	Data encrypted in transit using TLS 1.2 or higher, and at rest using AES-256.
Data retention and deletion	Program data retained for 60 days following program completion, then deleted. Earlier deletion available on request. Retention aligned to the client agreement.
Hosting and infrastructure	Independently audited cloud infrastructure (providers hold SOC 1/2/3 and ISO 27001). Segregated dev, staging, and production; secure SDLC with peer review and automated testing; centralized logging and monitoring; timely patching.
Subprocessors	Third-party vendors with data access pass a due diligence review and carry data protection and confidentiality clauses. Your coding tool and model provider are your own vendor relationships, not SI subprocessors, since SI is not in that path.

CONTROL AREA	WHAT WE DO
Incident response	An Incident Response Plan aligned to NIST SP 800-61, covering identification, investigation, containment, and remediation. You are notified of any confirmed incident affecting your data without undue delay.
Business continuity and backups	Encrypted backups stored in geographically redundant United States regions. Business continuity and disaster recovery plans tested and reviewed annually.
Personnel training	All personnel complete security and privacy training during onboarding and annually thereafter.
Legal and contractual	SI acts as a data processor under GDPR and CCPA, processing client data solely for the purposes defined in the client agreement.

5 Compliance and certifications

Superintelligent aligns its practices with ISO 27001 and the NIST Cybersecurity Framework and runs on certified cloud infrastructure. SI does not yet hold its own SOC 2 or ISO 27001 certification and maintains a roadmap toward SOC 2 Type 2 readiness. We state this plainly rather than imply a certification we do not hold.

ISO 27001	**Aligned**
NIST Cybersecurity Framework	**Aligned**
SOC 2 Type 2	**On roadmap, not yet certified**
Cloud infrastructure	**SOC 1/2/3 and ISO 27001 certified**

6 How the program improves your security posture

The program is security-positive by design. Rather than introducing agents without controls, it requires participants to build the governance around them as a condition of progressing. Among the artifacts your team produces:

ARTIFACT	WHAT IT ESTABLISHES
Security Manifesto	Per-agent security posture, credential management, permission scoping, and kill switches.
Governance Framework and Agent Evaluation Playbook	Tiered governance from personal-productivity agents through customer-facing ones, plus golden-set testing, stress testing, and production monitoring.
Third-party connector review	The practice of reviewing and allowlisting MCP servers and connectors, taught and applied during the program.

> The result is that your team leaves more capable of governing an agent workforce safely than when it started, with written standards your organization can apply beyond the cohort.

7 What we need from your security team to begin

Most prerequisites are self-serve. The items below are gated by IT, security, legal, or management and should be cleared before the cohort starts.

PREREQUISITE	OWNER
Approve the agentic coding tool (installed before day one)	Your security and IT
Approve Week 3 live connections (OAuth, admin consent)	Your IT or workspace admin
Allowlist the specific connectors and MCP servers	Your security
Decide document handling on the platform	Your security
Provide data-classification facts or an acceptable use policy	Your security or legal
Ring-fence participant time (about 5 to 7 hours per week)	Each participant's manager

The one operational gate is Week 3. Connecting agents to live systems depends on your environment being ready, so the people who can grant tool and data access should be reachable that week.

Contact

contact@besuper.ai

Please include "Agent Transformation Intensive" in the subject line.