GPTfy

Security · Privacy · Ethics · Compliance

Deploy AI Without Exposing New Attack Surface

GPTfy runs inside your Salesforce org and connects to **your** AI provider - Azure, AWS, or Google Cloud. No new vendors to validate. No new infrastructure to audit. **Your security team wraps up in days, not months.**

Days
Not months - to clear security review
Most customers complete within two weeks

Zero
External data centers
GPTfy managed package runs inside your org

4 Layers
Data masking before AI sees anything
Field, pattern, blocklist, Apex enforcement

Your security team has enough on their plate.

6+ Months

The queue never ends

Every team wants AI deployed yesterday. Every new AI tool means another round of vendor questionnaires, data flow maps, compliance certifications, and legal review. The backlog keeps piling up while the rest of the organization waits.

> *"There are quite a few third-party providers trying to do that, and all of them need me to send them accounts to a separate system. I'm not a big fan of that at all."*
> - CTO, FinServe Customer (East Coast)

Sprawling Surface

More vendors, more painstaking work

Every external API, data copy, and vendor server is another point to monitor. More risk registers. More pen tests. More incident response plans. Tedious, repetitive work that multiplies with every vendor your team has to vet.

> *"I don't want to send accounts to a separate system... This helps me to keep everything in Salesforce."*
> - CTO, Enterprise Debt Collection and FinServe

Shadow AI

Reviews drag on, ChatGPT doesn't wait

While reviews sit in the queue for months, your reps and agents are already using uncontrolled AI. The longer the backlog, the bigger the shadow AI risk becomes - and that creates even more work to clean up.

> *"We end up building stuff just for building it, and users just don't capture value out of it."*
> - VP of Technology, Financial Services

What if AI ran on infrastructure you've already secured?

TODAY
6-month review queues. External data centers to audit. Shadow AI spreading. Laborious vendor questionnaires piling up while the organization waits for AI.

→

TURN ON GPTFY
Install from AppExchange (Salesforce security-reviewed). Connect to YOUR Azure, AWS, or Google Cloud AI. Security & Trust Layer runs inside your org. Your admin controls every callout.

→

RESULT
Security review in days, not months - nothing new to validate. Zero new attack surface. Full audit trails. AI runs on infrastructure your team already secured.

THE S.P.E.C. FRAMEWORK - SECURITY · PRIVACY · ETHICS · COMPLIANCE

Enterprise AI governance in four layers

S Security: Mask Data Before AI Sees It

4-layer protection: field masking, regex patterns, blocklists, custom Apex. AI only processes masked data. AI firewall blocks prompt injection attacks.

+ Point-and-click config, zero code for Layers 1-3
+ Role-based masking (doctors vs. billing vs. compliance)
+ De-masking honors Salesforce field-level security
+ Prompt injection detection and prevention

P Privacy: Your Data. Your Infrastructure.

Complete data sovereignty. You choose where AI runs - US, EU, or APAC. Zero data retention with AI providers - data processed transiently, nothing stored.

+ GDPR, CCPA, CPRA compliant architecture
+ Regional data residency controls
+ Zero data retention with AI providers
+ Location-based AI access controls

E Ethics: Guardrails on Every Interaction

Prompt grounding rules enforce ethical, content, and security guardrails before AI responds. Four types of grounding - ethical, content, dynamic, and data security - applied to every prompt.

+ Bias detection and toxicity filtering
+ Hallucination prevention via content grounding
+ Dynamic grounding for locale and context
+ Human-in-the-loop review before action

C Compliance: Track and Prove Everything

Every AI interaction logged: user ID, timestamp, input, masked output, AI response, policies applied. Supervisor dashboards. E-discovery ready.

+ Fully traceable audit trails with field history tracking
+ Configurable retention and deletion policies*
+ One-click e-discovery export for regulatory exams
+ Real-time usage monitoring and supervisor dashboards

Your Infrastructure. Your AI. Your Control.

Bring your own model. Industry-specific compliance. Deploy in days, not months.

Zero-Trust Architecture - GPTfy's Security & Trust Layer Runs Inside Your Org

Everything runs in your controlled environment. Your IT team stays in control.

FRONTEND

Your M365 / Salesforce

Your SSO. Your conditional access policies. Salesforce Lightning, Copilot, Slack, Teams.

→

SECURITY & TRUST LAYER

Inside YOUR Salesforce Org

Managed package. 4-layer data masking, audit trails, ethics controls, and compliance enforcement. Admin-controlled callouts. Zero external servers.

→

AI BACKEND

YOUR AI Infrastructure

Azure OpenAI. AWS Bedrock. Google Vertex. Your contracts. Your data residency. Your security controls.

GPTfy cannot make external callouts unless your Salesforce Admin explicitly configures them via named credentials.

HOW WE BUILD - ENGINEERING RIGOR BEHIND EVERY RELEASE

430+	321+ hrs	Checkmarx	Code Escrow
Regression tests run per release	Testing investment per release	SAST security scan every AppExchange release	Source code held by Codekeeper for business continuity

BRING YOUR OWN MODEL

Use your AI infrastructure - not ours

No Vendor Lock-In

Switch from OpenAI to Claude to Bedrock without rebuilding. GPTfy adapts to your AI choices.

Your Contracts

Use existing Azure EA, AWS EDP, or GCP committed spend. Your negotiated pricing, not ours.

Data Residency

Choose Azure US East, AWS Frankfurt, GCP Singapore. Your compliance needs, your choice.

Full Transparency

Know exactly which model version processes your data. No opaque platform updates.

INDUSTRY COMPLIANCE

Pre-built configurations for regulated industries

Financial Services

FINRA/SEC compliance with MNPI protection and supervisor review dashboards.

- ✓ FINRA Rule 3110: AI interaction supervision
- ✓ Regulation S-P: PII/MNPI masking
- ✓ FINRA 4511: Configurable 6-year retention*
- ✓ SEC 17a-4: Tamper-evident record storage*
- ✓ One-click export for regulatory exams

Healthcare

All 18 PHI identifiers protected. BAA provided. With masking, your AI provider may never see PHI.

- ✓ §164.312(a): Access controls via SF profiles
- ✓ §164.312(b): Audit trails for every PHI access
- ✓ §164.312(e): AES-256 at rest, TLS 1.2+ in transit
- ✓ BAA simplification: masked data reduces scope
- ✓ Data never leaves your infrastructure

Insurance

NAIC Model Law #668 alignment. Claims processing with PII masked. 50-state breach readiness.

- ✓ Policyholder PII: SSN, DOB, DL, accounts
- ✓ Medical info safeguards in claims notes
- ✓ 50-state breach notification readiness
- ✓ Claims triage and fraud detection with masked data
- ✓ Configurable retention for state requirements

> *"There are quite a few third-party providers trying to do that, and all of them need me to send them accounts to a separate system. I'm not a big fan of that at all. This helps me to keep everything in Salesforce."*
>
> - CTO, Enterprise Debt Collection and FinServe Company

Questions From CISOs and Security Teams

DATA RESIDENCY

Does GPTfy store or process any of our data?

No. GPTfy is a managed package inside your Salesforce org. Your data remains in your infrastructure at all times - only a masked, sanitized version is transiently sent to your AI infrastructure for processing. Zero GPTfy servers. Zero caching. Zero data copies.

DATA MASKING

How does the Security & Trust Layer mask data?

Every field is reviewed against your masking rules before anything reaches AI. Four layers: (1) Field value - replace sensitive fields with tokens. (2) Regex - detect and mask PII across structured and unstructured data. (3) Blocklists - remove restricted terms that should never reach AI. (4) Custom Apex - your own masking, tokenization, or encryption logic. All configurable through point-and-click.

SALESFORCE SHIELD

Is GPTfy compatible with Salesforce Shield?

Yes. GPTfy runs as a managed package inside your org, so it inherits your Shield implementation. Platform Encryption, Event Monitoring, and Field Audit Trail all work as configured. If you've invested in Shield, that investment extends to your AI deployment with no additional configuration.

DATA RETENTION

Are hyperscaler data retention policies respected?

Fully. GPTfy calls your AI provider's API directly - Azure OpenAI, AWS Bedrock, Google Vertex - so all hyperscaler-level controls apply. Zero data retention configured on Azure? It applies to every GPTfy prompt. Content filtering, abuse monitoring, regional policies - all respected as per your infosec policies.

AI GOVERNANCE

Can we control which AI models users access?

Completely. Each prompt is configured once - with specific fields, data access rules, and an AI model - then assigned to one or more Salesforce user profiles. Your admin controls exactly who can run what, with which data, on which model. No duplication. No per-profile variants. Standard Salesforce profile management.

AI FIREWALL

Does GPTfy work with our AI firewall?

Yes. GPTfy's ConnectorClass architecture supports integration with internal or third-party AI firewalls. Authentication and processing classes route through your centralized AI security layer - your existing prompt injection detection, content filtering, and policy enforcement stay in the loop. Your security stays consolidated; GPTfy plugs into it, not around it.

MICROSOFT COPILOT

How secure is the Microsoft Copilot integration?

Three layers: (1) AppSource Marketplace security reviewed. (2) Runs in your Microsoft infrastructure - respects Entra ID, conditional access, and user permissions. (3) Azure connector between Copilot and Salesforce runs in your infrastructure. Code available for security review on request.

TRUST CENTER

What's in your security review packet?

AppExchange security approval documentation, Checkmarx SAST scan results, SOC 2-equivalent trust packet, shared responsibility matrix, data governance and access control policies, and incident response procedures. Available at gptfy.ai/trust-center

Review Our Security Posture

Trust packet, shared responsibility matrix, and compliance documentation.

Visit GPTfy Trust Center →

GPTfy

gptfy.ai

salesforce available on the AppExchange

Get it from Microsoft AppSource