BOLT BY SPARCLE · SOVEREIGN AI FOR THE ENTERPRISE

Governance, Data Sovereignty & Compliance

Put AI into every workflow, without sending your data anywhere you don't control and can't see.

Empower · Govern · Audit, for the people who have to answer for it.

“

Your people are already using AI on your data. The only question worth asking is whether that use is one you can control, one you can see, and one you can prove. Sparcle exists to make the answer yes.

WHY SPARCLE EXISTS

THE FORCING FUNCTION

AI adoption is outrunning AI control

The productivity is real and unstoppable. The governance to match it is not in place, and the liability still sits with you.

Shadow AI is already inside

Roughly half of employees use AI tools no one approved, and senior staff are the heaviest users. Every paste of a customer name or account number is an ungoverned data transfer.

Controls are missing where it matters

In breaches involving AI, **97%** of organizations lacked proper AI access controls, and **65%** of shadow-AI incidents exposed customer PII.

The accountability is yours

Regulators, boards and customers hold the organization responsible, not the model vendor. You are asked to prove what your AI touched, and where it went.

Sources: IBM Cost of a Data Breach 2025; Reco / CIO State of Shadow AI 2025.

THE MARKET'S BLIND SPOT

Most AI governance is documented, not enforced

The tools on the market split into two camps, and both leave the same gap open.

Paper, not a control

Policy packs, AI registries, model cards and risk registers describe intent. They sit beside the data flow, not in it. Nothing stops the next prompt.

Enforced, but outside your walls

Cloud proxies and cloud assistants that do enforce inspect your prompts inside **their** cloud, or only within one vendor's ecosystem. Your data still leaves the boundary you control.

The empty square

No dominant player enforces in the data path, inside your perimeter, in front of any model you choose. That gap is exactly where Bolt sits.

Framing: Gartner, "AI governance needs more than policies"; AI TRiSM runtime-enforcement layer.

WHY OUR MANTRA IS THREE VERBS

One operating model: Empower, Govern, Audit

Each verb protects the other two. Govern without Empower gets bypassed by shadow AI. Empower without Audit is indefensible. Together they are how responsible AI actually holds.

Empower

Give the whole workforce real AI leverage, across every system, so no one needs an unsanctioned tool. Most everyday work resolves on-device, at zero tokens.

Govern

Every use crosses a boundary you set: PII masked before the model, access scoped to identity, capabilities and models under admin control.

Audit

Every action is recorded in a tamper-evident log you can verify offline. You can answer "what did AI touch?" the moment you are asked.

01 · GOVERN · BOUNDARY

You define the boundary. Bolt enforces it.

Air-gap, a private model endpoint, or a public model with masking on. You pick the posture per deployment. Sparcle is never in the loop.

Runs in your infrastructure

Helm chart or Docker in your Kubernetes cluster, VPC, on-prem server, or air-gapped network. Lives entirely inside your perimeter.

- ALWAYS ON

Your LLM, your terms

Bring any model: a private endpoint on your GPUs, a local open model, or a public API. Admins can set an org privacy floor that refuses any model whose endpoint sits below it.

- ALWAYS ON
- PRIVACY FLOOR: ADMIN OPT-IN

Your keys, no lock-in

Bring your own keys, vaulted with per-tenant encryption. Swap models or providers without re-platforming.

- ALWAYS ON

GOVERN · MASKING · THE FLAGSHIP CONTROL

Sensitive data is masked before it reaches the model

One chokepoint on the LLM boundary, always on for structured identifiers, independent of which model you picked. Personal and health identifiers are detected, masked and pseudonymized before any prompt leaves your boundary, across streaming responses and text inside images, then restored on-device in the answer. Enforced in the path, not written in a policy.

1 · Identify

Detect personal and health identifiers in the prompt, contacts, financial and national IDs, validated with real check-digit algorithms so look-alikes are not flagged. Text inside attached images is scanned by OCR through the same detectors.

2 · Mask and pseudonymize

Replace each identifier with a stable, reversible pseudonym before the prompt leaves your boundary. The model reasons over the request, and on an unknown checksum the pipeline fails closed.

3 · Restore on device

Map the pseudonyms back to the real values inside your perimeter when the answer returns. The real values stay inside the boundary you control.

- ALWAYS ON
- SAME SAFETY FLOOR ON ANY MODEL, PUBLIC OR PRIVATE
- FREE-TEXT NAME DETECTION: ADMIN-ENABLED NER

GOVERN · ACCESS, CAPABILITY & OVERSIGHT

Policy and oversight, not trust

Who can use which systems, models and capabilities is set by your directory, and the consequential actions wait for a human. Nothing is left to the honor system.

Directory-driven RBAC

Access scoped by role, department and group from your IdP over SCIM or SAML. Every action reaches each system as the user's own identity, never a shared bot.

ADMIN OPT-IN

Capability lockdown

Enable, disable or restrict capabilities by role or group, connectors, models, clipboard, screen capture, with a kill-switch on each.

ADMIN OPT-IN

Versioned policy

Capture, simulate against your user population, activate, and roll back by pointer. Change the rules without fear.

ADMIN OPT-IN

Human approvals

Send, delete, schedule and pay wait for explicit human sign-off, a durable request that survives a restart, so nothing consequential auto-fires.

ALWAYS ON

GOVERN · AGENT GATEWAY · THE SHADOW-AI ANSWER

Govern the AI your people already use

The shadow-AI problem needs an answer, not just a policy. Bolt puts the traffic from the AI tools your people already reach for through the same masking, policy and audit boundary, so their convenience stops being your exposure.

A governed relay

Third-party AI clients route through Bolt. Prompts, tool calls and image content are masked, checked against policy, and logged before anything leaves your boundary.

ADMIN OPT-IN · CLAUDE TODAY

On-device AI-DLP

On paste, or as your people type into public AI sites, sensitive data is detected and redacted in place, and can be blocked outright under a managed policy.

WARN BY DEFAULT · BLOCK BY POLICY

The same evidence

Every mask, allow and deny on that traffic seals into the same audit chain. Shadow AI becomes governed AI you can prove.

ALWAYS ON

02 · COMPLY · CONTROL MAPPING

The same controls, mapped to every framework

Buyers now expect this mapping as table stakes. The differentiator is that on Bolt these are running controls, not documented intentions.

FRAMEWORK	WHAT IT ASKS FOR	THE BOLT CONTROL
NIST AI RMF + GenAI Profile	Privacy-enhanced measurement, third-party risk management, record-keeping	In-path PII masking, model catalog governance, audit on every request
EU AI Act	Logging and record-keeping, human oversight, data governance	Merkle-sealed logs, durable approvals, boundary control and masking
ISO/IEC 42001	AI management controls: data for AI, responsible use, third-party	Policy engine, capability governance, evidence-pack export
ISO 27001 / SOC 2	Access control, logging, cryptography, data-leak prevention	RBAC, per-tenant encryption, tamper-evident audit, masking as DLP
CSA STAR (CAIQ v4)	Cloud-controls self-assessment across the CCM control domains	Published CAIQ v4 answer set and evidence map, self-hosted deployment model
India DPDP (Rule 6)	Encryption or masking, access logging, logs kept one year, breach reporting	Masking, per-tenant CMK, retention sweepers, sealed audit trail
GDPR / CCPA	Data minimisation, DSAR and erasure, pseudonymisation, opt-out of automated decisions	Erasure cascade, pseudonymisation at the boundary, human-in-the-loop

COMPLY · REGULATION PACKS

One engine, any jurisdiction

Regulations ship as signed, swappable packs that feed the same masking chokepoint. A new country is one pack, not a new product.

Built-in coverage

Packs across major jurisdictions, GDPR, UK GDPR, DPDP, CCPA, LGPD, PIPL, APPI, PIPA and more, each declaring its identifiers, rights and residency.

ALWAYS ON

Real identifier validation

National IDs are checked with their true check-digit algorithms (Aadhaar, PAN, GSTIN, NHS, Singapore NRIC, Australia TFN, Brazil CPF and CNPJ, Japan MyNumber, Korea RRN and more), so look-alike strings do not trigger false masking.

ALWAYS ON

Signed and additive

Packs are cryptographically signed bundles, reviewed on install, activated per organization. Author a new jurisdiction without touching the engine.

ALWAYS ON

COMPLY · INDIA DPDP ACT

Govern the point of use, the leg no consent manager can see

Consent is captured at your front door by your consent manager. Bolt sits downstream, governing the internal leg where your people and your AI actually touch that data.

THE BLIND SPOT TODAY

The internal AI leg

- ✗ Consent managers and CRMs record that consent exists
- ✗ They are not in the path when an employee runs AI on the data
- ✗ No CMP, DLP or CRM can prove that leg honored the purpose

WHAT BOLT ADDS

Enforcement plus proof, in your perimeter

- ✓ Purpose enforced at the point of use, derived from context and policy
- ✓ PII masked before any model, on-prem or BYO-LLM
- ✓ Signed record that your people and AI honored consent

● MASKING, AUDIT, RESIDENCY: ALWAYS ON

● CONSENT-LINKAGE PURPOSE ENFORCEMENT: ADMIN OPT-IN

● LIVE CMP VENDOR WIRING: IN PROGRESS

DPDP Rules notified Nov 2025; substantive obligations from May 2027. Position Bolt as your data processor, downstream of your consent manager.

COMPLY · DATA SUBJECT RIGHTS

Honor the rights, on a clock you keep

Erasure, retention and key destruction are wired as running jobs, mapped to the rights each active regulation declares.

Erasure and DSAR

A typed cascade erases a subject across messages, sessions and pseudonym tokens, tenant-scoped so one erasure never crosses into another tenant, and revokes the user's upstream OAuth grants so a deleted user keeps no live access.

ALWAYS ON

Retention sweepers

Messages, sessions, runs and audit entries age out on the retention window you set, enforced daily, not left to manual cleanup.

ALWAYS ON

Crypto-shred with receipt

Destroy a tenant's key and its data becomes unrecoverable, the strongest form of erasure. Each erasure emits a signed, offline-verifiable deletion receipt sealed into the audit chain when KMS is configured.

ALWAYS ON

03 · DATA SOVEREIGNTY · TOPOLOGY

Where your data lives

The workforce, the engine, the masking, and the store all sit inside a boundary you own. Only masked prompts cross to the model you selected. Sparcle has no server in the path.

Sparcle is never in the data path. There is no hosted route back to us.

77% of buyers weigh a vendor's country of origin in AI purchasing decisions. Source: Deloitte 2025.

SOVEREIGNTY · ENCRYPTION

Encryption and key control, held by you

Data at rest is encrypted under per-tenant keys you control, so possession of the disk is not possession of the data.

Per-tenant keys, your KMS

Envelope encryption with a customer master key per tenant, backed by a local KMS, AWS KMS or HashiCorp Vault. Bring your own keys.

ALWAYS ON

Token vault

OAuth and model credentials sealed with authenticated encryption, bound so a token cannot be lifted and replayed for another user.

ALWAYS ON

Crypto-shred and rotation

Rotate keys on schedule; shred a tenant key to make its data unrecoverable. Each shred emits a signed, offline-verifiable deletion receipt, so erasure is proven, not just asserted.

ALWAYS ON

SOVEREIGNTY · THE SAFETY NET

Even the most permissive choice stays safe

The question every CISO asks: "if my analyst picks a public model, does sensitive data leak?" The answer is no, because masking happens before the call, regardless of the model.

CONSUMER AI

Convenience, no floor

- ✕ Raw PII typed straight into the model
- ✕ No record of the transfer, no way to mask
- ✕ Data retained and processed outside your reach

BOLT, ANY MODEL

Convenience with a floor

- ✓ Identifiers pseudonymized before the prompt leaves
- ✓ The model reasons over tokens, never the real values
- ✓ Real values restored on-device, and the event is logged

● A PRIVACY FLOOR NO CONSUMER AI OR COPILOT CAN MATCH

04 · AUDIT · THE PROOF

A record you can verify, without trusting us

Most tools give you a log you have to take on faith. Bolt gives you a cryptographic proof of integrity that a regulator can check with a standalone verifier, needing neither our database nor our keys.

Tamper-evident chain

Events are hashed into a Merkle tree and the root is signed by your key. Altering any past entry breaks the proof.

ALWAYS ON

Offline-verifiable export

Export an evidence bundle that verifies on its own. No live connection to Sparcle, no dependence on our infrastructure to prove it.

ALWAYS ON

The differentiator

Competitors hand you a log. Bolt hands you proof. For a compliance buyer, that is the difference between an assertion and evidence.

ALWAYS ON

AUDIT · ACCOUNTABILITY BY DESIGN

Accountability at every layer

Governance is not one gate. It is a stack of controls, each enforcing something and each emitting evidence, so no path can skip a rule and every question has an answer.

- **Identity** Every action runs as the user's own identity, via your IdP and SCIM.
- **Boundary** Data moves only to endpoints you designate. Air-gap capable.
- **Masking** PII and PHI pseudonymized before any model call, on any model.
- **Policy** RBAC and per-capability rules from your directory attributes.
- **Approval** Durable human sign-off on destructive or outward-facing actions.
- **Audit** A Merkle-sealed, offline-verifiable record of all of the above.

AUDIT · TRANSPARENCY

The evidence you hand an auditor or a board

Transparency is not a dashboard you hope is right. It is a record you can produce, filtered and signed, the moment you are asked.

Every request logged

Method, path, user, status, duration and correlation IDs on every API call, with credential paths redacted.

ALWAYS ON

Evidence pack by regulation

Export the safeguards and audit history for a specific regulation, labeled and ready for an assessment or inquiry.

ALWAYS ON

Answer the hard question

"What did our AI touch, for whom, and where did it go?" is a query, not a research project.

ALWAYS ON

Know your coverage

An org-wide read of whether masking and policy cover every surface, and whether any config has gone stale, so assurance is measured, not assumed.

ALWAYS ON

THE HONEST POSTURE

What you can attest to today

We don't sell you a certification. We give you the running controls that get you certified, and we're precise about what's always on, what your admin turns on, and what's still roadmap.

ALWAYS ON

- PII / PHI masking, text and images
- Merkle-sealed, offline-verifiable audit
- Per-tenant encryption, BYOK, shred
- Signed deletion receipts, OAuth revoke
- Regulation packs with real ID checksums
- Durable human approvals
- BYO-LLM, in-perimeter deployment

ADMIN OPT-IN

- Directory-driven RBAC
- Capability lockdown by role and group
- Versioned policy: simulate and rollback
- Command and terminal governance
- Agent Gateway for third-party AI (Claude)
- Model privacy-floor enforcement
- In-app egress allow-listing
- DPDP consent-linkage enforcement

DESIGN PARTNER ROADMAP

- DSR intake console with SLA clocks
- Live consent-manager vendor wiring
- Independent read-only auditor role

THE ASK · EMPOWER · GOVERN · AUDIT

AI you can put your name behind.

Enforced, in your perimeter, in front of any model. Empower the whole workforce, govern every use at a boundary you control, and audit it with proof you can hand over, without surrendering control or visibility to a cloud provider.

3
VERBS, ONE OPERATING MODEL

0
DATA PATHS BACK TO SPARCLE

bolt@sparcle.app
START THE CONVERSATION