

Bolt + Aeira · Enterprise AI Substrate Architecture

Inside your perimeter. With your LLM. On your audit trail.

- Customer network boundary · Bolt Absolute & Bundled tiers · No outbound to Sparcle · No telemetry -

① **BYO LLM here** — customer's existing endpoint, any vendor.

② **Audit chain seals here** — 10-second Merkle epochs, KMS-signed.

③ **ACL pre-filter here** — 4-axis GIN index at query time.

④ **Webhook receivers here** — push-mode for top-10 sources **PILOT**

⑤ **Cross-source action verbs via MCP** — Bolt invokes live source APIs.

⑥ **Connector poll here** — per-source cadence, 5 min to hourly.

⑦ **Per-region keyrings** — sovereignty enforced at write + read **PILOT**

⑧ **Permission cache** — L1 LRU shipped; L2 Redis **PILOT**

⑨ **Late-bind check** — top-K only, 6 connectors shipped today.

Bolt Desktop + PWA + Browser Extension

Cmd-Space launcher Native desktop, code-signed	Web app installable on desktop + mobile browsers	Browser extension cmd palette anywhere	Local bridge ws://127.0.0.1, WebCrypto pairing
--	--	--	--

Persona-aware utility packs · Scoped action verbs · Cross-source workflows via MCP · Voice + keyboard · Append-only signed session history

Aeira ACL-Aware Search + Audit Substrate

SEARCH PIPELINE

API server axum · JWT	Permission classifier Flatten / Tuple / Hybrid	ACL pre-filter 4-axis GIN platform-locked	Keyword + Vector FTS · pgvector	RRF fusion top-K bounded	Authority Policy SDK your code · Rust + YAML customer-extensible	Late-bind top-K only	Audit emitter signed events
---------------------------------	--	---	---	------------------------------------	--	--------------------------------	---------------------------------------

STORAGE + CRYPTOGRAPHY

Per-tenant CMK Vault · AWS KMS · HSM	At-rest encryption per-tenant keys	Crypto-shred GDPR Art. 17 · HIPAA erasure	Merkle audit chain Ed25519, 10s seals	Standalone verifier customer auditor runs
--	--	---	---	---

INGEST SPINE + CACHE

NATS spine per-tenant subjects	Ingest saga worker PILOT	PII detection mask at LLM boundary	Permission cache L1 LRU shipped, L2 Redis PILOT	Webhook receiver PILOT
--	---	--	--	---

Connectors 103 in-house + Airbyte bridge (~350 sources)

Productivity Jira · Confluence · Slack · Teams · GitHub · GitLab · Linear · Notion · Box · Dropbox · GDrive · OneDrive · SharePoint	Mail + calendar Gmail · Outlook · GCalendar · Exchange archive	CRM + sales Salesforce · HubSpot · Outreach · Salesloft · Marketo	ERP + finance SAP S/4HANA · Oracle EBS · NetSuite · Workday Financial · Concur · Coupa · Anaplan · Workiva
HR + HCM Workday · BambooHR · Greenhouse · Lattice · SuccessFactors	Data + BI Snowflake · Databricks · BigQuery · Redshift · MongoDB Atlas · Tableau · Looker · Power BI	Support + ITSM Zendesk · Freshdesk · Freshservice · ServiceNow · Intercom · HelpScout	ECM + WCMS FileNet · Documentum · Alfresco · AEM · Sitecore · Drupal · WordPress · Domino · Quip
Security + observability CrowdStrike · SentinelOne · Tenable · Qualys · Vanta · Drata · Proofpoint · Mimecast · Veritas Vault · Splunk · Datadog · NewRelic · Sentry · PagerDuty	IDP + IT + iPaaS + CDN Okta · Entra · OneLogin · Jamf · MuleSoft · Boomi · Cloudflare · Akamai · Fastly	e-sign + collab + analytics DocuSign · Adobe Sign · Airtable · Smartsheet · Monday · Asana · ClickUp · Wrike · Trello · Zoom · Mixpanel · Pendo	Generic adapters + External REST · JDBC · WebCrawler (config-driven) + ~350 via Airbyte container bridge

Three tiers: Deep (~25) source-fidelity ACL + per-source enrichment · Standard (~80) coarser ACL · Airbyte bridge (~350) long-tail. 4-axis ACL emission · IR with Flatten/Tuple/Hybrid materialization · Late-bind on 6 sources today (Filesystem · Outlook · ServiceNow · Splunk · OneDrive · Salesforce); backfilling per pilot. **Enricher chain** at ingest pulls source-specific metadata (Confluence status · FHIR evidence · regulatory tier) into the Authority Policy signals.

★ **AUTHORITY POLICY SDK**
The customer-extensible wedge. Your domain's notion of "authoritative" lives in your code, in your perimeter. Reference impls + vertical packs.

↔ **Customer IDP**
Okta · Entra · OneLogin · SAML · OIDC

→ **Customer LLM**
OpenAI · Anthropic · Bedrock · Azure OpenAI · on-prem Llama / Mistral / Qwen

→ **Customer KMS**
Vault · AWS KMS · PKCS#11 HSM

TECH STACK

language	Rust 1.85+
workspace	~130 crates
api	axum · tower
runtime	tokio
search	Postgres FTS + pgvector
fusion	RRF (weighted)
policy	Authority Policy SDK
spine	NATS JetStream
cache L1	LRU
cache L2	Redis [pilot]
rebac	OpenFGA [pilot]
storage	FS + S3 Object Lock [pilot]
kms	Vault · AWS · PKCS#11
auth	OIDC · SAML · JWT
audit	Ed25519 + Merkle
telemetry	OpenTelemetry · Prom
desktop	Native · macOS · Linux · Windows
web	Svelte 5 · installable
browser ext	Chrome Manifest v3
mobile	iOS · Android [Q3 2026+]

Bolt tier guide · Pick by buyer type

Absolute \$30/seat · customer self-hosted · BYO LLM mandatory · **regulated & on-prem** · zero Sparcle sub-processors. **Bundled** \$60/seat · customer self-hosted + managed AI router or BYO LLM · 10M tokens/mo · **regulated mid-market** when BYO LLM is preserved. **Complete** \$90/seat · Sparcle-managed hosting · 15M tokens/mo · 99.9% SLA · **not for regulated workloads**.

IDP coverage · Authentication-side

Native: Okta · Microsoft Entra · OneLogin · Auth0 · Google Workspace · Zoho. **Via generic OIDC + SAML 2.0** (any compliant IdP): Ping Identity · Keycloak · AWS Cognito · AWS IAM Identity Center · ForgeRock · JumpCloud. **Directory connectors** ship for Okta, Entra, OneLogin; Ping & LDAP/AD on pilot-deliverable list.

✓ What's shipped today

- 103 in-house source connectors + Airbyte bridge (~350 sources)
- 4-axis ACL pre-filter at index time (provable in audit chain)
- Authority Policy SDK + 3 reference impls (Default · Regulated · Healthcare)
- Sealed Merkle audit chain + standalone verifier binary
- 3 deployment SKUs: trial sidecar · docker-compose · Helm
- 4 KMS variants: Vault · AWS KMS · PKCS#11 HSM · local-dev
- Cross-product audit byte-compatible between Bolt + Aeira

🔔 What ships during pilot

- End-to-end ingest saga (1-2 days from green-light)
- Real-time webhook receivers for top-10 sources
- L2 Redis permission cache + webhook invalidation
- OpenFGA wire-up for tuple-mode ACL
- Late-bind on next 14 connectors (per design partner priority)
- First published latency benchmark (target p95 ≤ 200ms on 100M-doc tenant)