

The Zero Trust Handbook for the Agentic Business

By ChordianAI

Foreword

Every company is about to gain a new kind of workforce. Agents will research, monitor, draft, check and decide, around the clock and at a scale no team could match. The question is no longer whether to deploy them. It's what they will know, and who controls it.

This handbook is for the people who have to answer that question: CEOs, CTOs, CISOs and heads of data. It's practical by design. Each chapter ends with decisions to make or a short checklist.

Chapter 1: The agentic shift

For thirty years, enterprise software was built for people clicking through screens. Agents don't click. They call APIs, read thousands of pages, and act on what they find. When agents become the main consumers of data, three things change.

Volume. An agent can run hundreds of searches before a person finishes reading one report. Data infrastructure built for occasional human queries breaks under continuous machine demand.

Speed of error. A person who reads a wrong number usually hesitates. An agent acts on it, and passes it to the next agent. Bad data no longer slows decisions down; it multiplies wrong ones.

Exposure. Every agent is a new identity with access to company knowledge. Without controls, each one is a new path for intellectual property to leave the building.

The agentic company is therefore not a model problem. Models are increasingly interchangeable. It's an intelligence problem: what agents know, how current and correct it is, and who controls it.

Decisions to make

- Which workflows will agents run first?
 - Which data will those agents need, and where does it live today?
 - Who owns the risk when an agent acts on wrong or leaked information?
-

Chapter 2: Why agentic companies need a context layer

A model is only as good as the context it receives. Give it stale, conflicting or unsourced data and it will produce confident, wrong answers. Give it everything and it will leak what it shouldn't.

A **context layer** sits between your data and your agents. It does four jobs:

1. **Collects** intelligence from the live web and your own systems
2. **Cleans and resolves** it: duplicates merged, entities unified, conflicts surfaced
3. **Maintains** it, so facts stay current and stale ones lose confidence
4. **Governs** it, so every agent sees only what it's allowed to see, and every access is recorded

Without a context layer, every team builds its own pipelines, every agent re-reads the same documents, and nobody can say where an answer came from. With one, every agent starts from the same governed source of truth.

What to look for

- Sources on every finding
 - Visible contradictions, never silently resolved
 - One identity per person, company and deal across systems
 - Permission scoping per user and per agent
 - Deployment where your data has to live
-

Chapter 3: Live data intelligence

Most of the information agents need isn't inside your company. It's on the open web, in registries, filings, regional sources and specialist sites. Static databases bought once a year can't keep up; they decay by design.

Live data intelligence means collecting at the moment of need, at machine scale, and turning raw data into something agents can trust. In practice:

- **Breadth:** many sources, many countries, many languages. Seeing a market the way a local sees it requires searching from that location.
- **Depth:** multi-hop research that plans, re-queries and cross-checks, not a single search.
- **Structure:** extraction into clean, typed data, not pages of text.
- **Cleaning:** duplicates removed, entities resolved, conflicts flagged with both sources visible.
- **Continuity:** scheduled searches that notice change and alert the right people and agents.

Measure it by

- **Coverage:** sources and countries reached per question
 - **Accuracy:** share of findings with verifiable sources
 - **Cleanliness:** duplicates removed and conflicts caught per report
 - **Freshness:** age of the newest supporting source
-

Chapter 4: Knowledge Twins

Every organization has desks whose knowledge is its edge: the investor-relations team that knows every mandate, the diligence team with its own checklists, the research team with its methodology. That knowledge is fragile. It lives in people, and people move on.

A **Knowledge Twin** packages a desk's intelligence and know-how into a capability agents and people can use around the clock. It combines:

1. **Live data intelligence:** what the world says today
2. **Institutional knowledge:** what the desk knows and has decided
3. **The desk's way of working:** trusted sources, standard checks, the questions it always asks

How to scope your first Twin

- Pick a desk with high-value, repeatable questions
- List the ten questions it answers most often
- List the sources it trusts and the checks it always runs
- Define who, and which agents, may ask what

How to measure it

- Questions answered without escalation
 - Time from question to sourced answer
 - Knowledge retained when team members change
 - Errors caught by contradiction flags before decisions
-

Chapter 5: Zero Trust for agents

Zero Trust started as a network principle: never trust, always verify. For agents it has to go further, because agents aren't just users of the network. They are users of knowledge.

The agentic threat model

- **Over-broad access:** an agent connected to everything can reveal anything

- **Prompt-driven leakage:** a clever request extracts what a policy should have blocked
- **Model egress:** sensitive context sent to external models
- **Untraceable actions:** nobody can reconstruct what an agent saw and why it acted

The controls that answer it

- **Identity for every agent,** not just every person
- **Attribute-based access control (ABAC):** results scoped to the requester, the data's sensitivity and the context
- **Tenant isolation,** tested continuously
- **Answer-time DLP:** sensitive data discovered and masked before it reaches a response
- **BYOK:** the customer holds the encryption keys
- **Tamper-evident audit:** every access, query and agent action recorded
- **Sovereign deployment options:** your cloud, on-prem or air-gapped

Measure it by

- Permission-leak rate (target: zero, contract-tested)
 - Share of sensitive items masked at answer time
 - Share of agent actions fully reconstructable from the audit trail
 - IP exposure incidents prevented
-

Chapter 6: Sovereign AI

Sovereignty is a deployment decision, not a geography. It answers one question: who controls where intelligence lives, which model processes it, and who can see it?

The four levels

1. **Managed cloud:** fastest to start; isolation per tenant
2. **Your cloud (BYOC):** data never leaves your VPC; you hold the keys
3. **On-prem:** your servers, your perimeter
4. **Air-gapped:** no external connectivity; answers generated by a sovereign model on your own infrastructure

The trade-off, measured Running fully sovereign has a cost in raw accuracy, and that cost should be published, not hidden. On the public LoCoMo benchmark (1,540 questions), Chordian scores 71.4% fully sovereign with the Chordian Sovereign Model and 81.3% with governed frontier synthesis on identical retrieval. The right choice depends on the workload. Many organizations run both: sovereign for sensitive work, frontier where policy allows.

Decisions to make

- Which workloads must never leave your infrastructure?
 - Which may use external models, under which policy?
 - Who approves exceptions, and how are they logged?
-

Chapter 7: Measuring value

Time saved is the smallest part of the return. Measure the full value of agentic intelligence across six dimensions:

Dimension	Example metric
Data volume	Sources and countries covered per question
Data accuracy	Share of findings with verifiable sources
Data cleaning	Duplicates removed, conflicts caught
Intelligence	Decisions informed; signals caught before competitors
IP protected	Sensitive items masked; exposure incidents prevented
Efficiency	Analyst hours, external spend and tokens per answer reduced

Report these monthly from the first pilot. They make the case for scaling better than any projection.

Chapter 8: The 90-day rollout

Days 1-30: one desk, one Twin

- Choose a desk with repeatable, high-value questions
- Connect its sources and documents
- Define access policies and audit requirements
- Run the desk’s ten most common questions; compare against its current process

Days 31-60: prove and extend

- Add Scheduled Search for the signals the desk watches
- Report the six value dimensions

- Extend to two more desks with similar needs

Days 61–90: agents at scale

- Give approved agents access via MCP or API, with per-agent identities
 - Review the audit trail with security and compliance
 - Decide the long-term deployment model: cloud, BYOC, on-prem or air-gapped
 - Set the roadmap for the next five Twins
-

Appendix A: Zero Trust checklist for AI vendors

- Is every finding returned with its source?
- Are contradictions shown, not silently resolved?
- Is access scoped per user and per agent?
- Is tenant isolation tested, and is the result published?
- Is sensitive data masked at answer time?
- Do we hold the encryption keys?
- Can it run in our cloud, on-prem or air-gapped?
- Is every agent action in a tamper-evident log?
- Is our data ever used for training? (The answer must be no.)
- Does the vendor publish benchmarks, including weak results?

Appendix B: Glossary

- **ABAC:** attribute-based access control; access decided by attributes of the requester, the data and the context
- **BYOK:** bring your own key; the customer controls encryption keys
- **Context layer:** the governed layer between data and agents
- **Entity Coherence:** one consistent identity for every person, company and deal across systems
- **Knowledge Twin:** a desk's live intelligence and institutional knowledge, packaged for agents and people
- **Sovereign model:** a model that runs entirely on the customer's infrastructure

Appendix C: Policy starter templates

1. **Agent access policy:** which agents, which data, which actions
2. **External model policy:** when sensitive context may leave the perimeter

3. **Audit review policy:** who reviews agent logs, how often, and what triggers escalation