SALT

AI & Data Strategy Whitepaper

Practical framework for implementing AI and data-driven strategies in your organization

Salt Technologies, Inc.
Software Consulting & Development Partner
Austin, TX, USA | Pune, MH, India

1. Why you need an AI & data strategy now

Every leadership team is being asked *“What is our AI strategy?”*

At the same time, many organizations are stuck in proof-of-concept purgatory: scattered pilots, demos that never reach production, and dashboards no one uses.

An AI & data strategy is not a slide deck about “using AI everywhere.”

It is a practical plan that connects business outcomes, data foundations, and delivery capabilities into a repeatable system.

In this whitepaper, you’ll get:

- A simple framework for connecting AI initiatives to measurable business value.
- A practical view of MLOps – the engine that keeps AI models reliable in production.
- Concrete patterns for using large language models (LLMs) in real workflows.
- A checklist for responsible AI governance so you can move fast without losing control.

You can use this as a blueprint for your own roadmap or as a reference when evaluating partners.

2. What a good AI & data strategy actually looks like

A useful AI & data strategy answers four questions clearly:

1) Where will AI and data move the needle for our business?

- **Revenue:** higher conversion, cross-sell/upsell, reduced churn.
- **Cost:** process automation, better capacity planning, fewer errors and rework.
- **Risk:** fraud detection, compliance monitoring, early-warning signals.
- **Experience:** faster response times, personalization, better employee tools.

2) What data and technical foundations do we need?

- A clear view of what data exists today, where it lives, and its quality.
- A modern data architecture that can support analytics, ML, and LLMs.
- Secure, governed access so the right people and systems can use the data.

3) How will we deliver AI safely and repeatably?

- Ways of working: product mindset, cross-functional teams, experimentation.
- MLOps practices: versioning, testing, deployment, monitoring for models.
- LLM application patterns: how to plug generative AI into real processes.

4) How will we manage risk and governance?

- Policies on privacy, security, IP, and acceptable use.
- Guardrails for model behavior and human-in-the-loop review.
- Clear ownership for AI systems once they're in production.

If your current AI efforts don't address these four questions, you probably have isolated experiments rather than a real strategy.

3. Foundations: preparing your organization for AI

3.1 Start from business outcomes, not algorithms

The most successful AI programs start with a small set of high-value use cases, not a shopping list of tools.

Practical steps:

1. **Map your value drivers**
   Identify 3–5 key metrics you care about (for example, average handling time, on-time delivery, repeat purchase rate, claims processing time).

2. **Walk the process**
   Sit with operations, customer support, finance, sales – whoever owns those metrics. Look at where delays, handoffs, manual checks, and rework happen.

3. **Capture opportunities**
   For each pain point, ask:
   "Could better data, prediction, or automation significantly improve this step?"

Your output should be a simple **use-case backlog**:

- Problem statement
- Impact metric (and a rough upside estimate)
- Data needed
- Effort/complexity

From there, you prioritize on impact vs. feasibility instead of chasing trendy AI use cases.

3.2 Build a realistic data foundation

You don't need a perfect, futuristic data platform to start. But you do need some basics.

Key elements:

- Source systems inventory
  - List your core systems (ERP, CRM, product, support, finance, marketing, etc.).
  - Note what data they hold, how often it changes, and who owns it.
- Data integration and storage
  - Decide where integrated data will live (data warehouse, lakehouse, or similar).
  - Use repeatable pipelines (ELT/ETL) instead of manual exports and spreadsheets.
- Data quality and governance

    - Identify “golden sources” for key entities (customer, product, order, asset).
    - Define a few simple quality rules: completeness, uniqueness, valid values.
    - Establish data ownership: who is responsible for fixing issues in each domain.
- Security and access
    - Role-based access rather than one-off approvals.
    - Audit trails for who accessed what.
    - Clear policies for using production data in development and testing.

Think of this as a **minimum viable data platform** that can grow as your AI usage grows.

3.3 Operating model and skills

AI and data projects fail when they are “thrown over the wall” between teams.

You will need:

- Cross-functional delivery pods that include:
    - Product owner (business side, owns outcomes)
    - Data/ML engineers and software engineers
    - Data scientist / ML engineer where needed
    - Domain experts from operations or customer-facing teams

- Enabling functions:
    - Data platform team that provides shared tooling
    - Security, compliance, and architecture for guidelines and guardrails
    - Change management and training to drive adoption

You don’t have to build all of this on day one, but you do need clarity on who owns what.

4. A practical AI & data strategy framework

Here is a simple four-step framework you can apply in your organization.

Step 1: Discover

- Identify high-value, data-driven opportunities.
- Prioritize a small number of use cases based on impact and feasibility.
- Clarify success metrics and constraints for each use case.

Deliverables:

- Use-case backlog with impact/effort scores.
- Baseline metrics for top 3 use cases.

Step 2: Design

- Design the data flows, model approach, and integration points.
- Decide where in the process AI will act (assist, recommend, or automate).
- Define the user experience: who will consume the insights and how.

Deliverables:

- Solution sketch and architecture at a high level.
- Data requirements and data quality assumptions.
- Governance and risk assessment for the use case.

Step 3: Deliver (Pilot with MLOps/LLMOps from day one)

- Build an end-to-end pilot that runs on real data and real workflows.
- Apply MLOps practices:
    - Version control for code and models
    - Automated tests and data validation
    - CI/CD pipelines to deploy models and application code
    - Monitoring for data drift, model performance, and system health

Deliverables:

- Working pilot in production for a limited audience or scope.
- Initial performance results and feedback from users.

Step 4: Govern & Scale

- Decide whether to scale the pilot based on clear criteria (business impact, reliability, risk).
- Harden the solution: resiliency, documentation, runbooks, and handover.
- Integrate into your broader AI governance process.

Deliverables:

- Production-grade AI service or application.
- Playbook for operating and improving the solution.
- Lessons learned that feed into the next cycle.

This framework is repeatable: every new use case runs through **Discover → Design → Deliver → Govern**, getting faster and safer over time.

5. MLOps: keeping AI models healthy in the real world

Without MLOps, AI systems degrade silently. Data drifts, user behavior changes, and models that once performed well become risky.

5.1 The typical ML lifecycle

- Problem framing and data exploration
- Feature engineering and model training

- Evaluation and selection between models
- Deployment into production (batch, streaming, or real-time)
- Monitoring, feedback collection, and retraining

5.2 Core MLOps capabilities you actually need

You don't need every tool on the market. Focus on these foundations:

- Source control and environments
  - All code (including experiments) lives in version control.
  - Reproducible environments using containers and infrastructure-as-code.
- Experiment tracking
  - Log model configurations, metrics, and data versions used for each run.
  - Make it easy to compare experiments and roll back if needed.
- Model packaging and deployment
  - Standard way of serving models (e.g., APIs or batch jobs).
  - Automated deployment pipelines with approvals where needed.
- Monitoring and alerts
  - Track technical metrics (latency, error rate) and ML metrics (accuracy, drift).

    - Alert when performance moves outside defined thresholds.
    - Define clear runbooks for responding to issues.
- Feedback loops
    - Capture user feedback and ground-truth outcomes.
    - Use that feedback to retrain or adjust the model in a controlled way.

5.3 A simple example

Imagine a **demand-forecasting model** for a retail chain. With basic MLOps in place, you would:

- Validate incoming sales data before training and scoring.
- Track which model version is generating each forecast.
- Monitor forecast error over time and by category or region.
- Trigger retraining when data drift or performance issues are detected.
- Provide alerts to the operations team if the model is outside safe ranges.

The result is a model that stays useful instead of becoming a black box.

6. LLM applications: where large language models really help

Large language models (LLMs) are powerful, but using them safely in your business requires structure.

6.1 Common high-value patterns

Pattern 1: Knowledge assistant for employees

- Use case: Answer internal questions about policies, processes, and product information.
- Approach: Combine an LLM with retrieval from your own documents (RAG – retrieval-augmented generation).
- Benefits: Less time searching for information, fewer repetitive questions to subject-matter experts.

Pattern 2: Workflow copilot

- Use case: Help agents or internal users draft responses, summarize cases, or fill forms.
- Approach: LLM suggests content, humans approve or edit, systems log final output.
- Benefits: Faster handling times, more consistent communication, reduced burnout.

Pattern 3: Classification and extraction

- Use case: Categorize tickets, extract key fields from documents, detect intent in messages.

- Approach: Use an LLM (or smaller models) to label or extract structured data from unstructured text.
- Benefits: Better routing, improved reporting, cleaner data for downstream models.

6.2 Design principles for LLM applications

- Stay close to real workflows
  - Embed the assistant or copilot in the tools people already use.
  - Design for assistive use first, full automation later.
- Ground outputs in your data
  - Use retrieval from trusted internal sources instead of asking the model to "make things up."
  - Show citations or links back to the original documents where possible.
- Keep a human in the loop
  - Make it easy to review and correct model outputs.
  - Decide clearly who is responsible for the final decision or message.
- Track usage and impact
  - Measure productivity (time saved), quality (error rates), and satisfaction.
  - Use these metrics to decide where to invest further.

6.3 LLMOps considerations

For LLM-based systems, you should also:

- Version prompts, configuration, and model choices.
- Monitor cost, latency, and failure modes (e.g., rate limits).
- Define clear fallbacks when the LLM cannot answer reliably.

7. Responsible AI and governance

Moving fast with AI is only sustainable if you have clear guardrails.

Key risk areas to manage:

- Privacy and security
    - Protect personal and sensitive data at rest and in transit.
    - Avoid sending confidential data to external services without proper controls.
- Bias and fairness
    - Understand where training data and labels come from.
    - Regularly test for unwanted bias in model outputs, especially in high-impact decisions.
- Reliability and explainability
    - Make sure critical decisions are auditable.
    - Provide explanations or supporting evidence where required by regulators or customers.
- Intellectual property and content

    - Decide what content your organization is comfortable generating with AI.
    - Clarify ownership and licensing of generated content.
- Compliance and regulation
    - Map AI use cases to applicable regulations in your geography and industry.
    - Maintain documentation of models, data sources, and evaluations.

Governance in practice looks like:

- An AI steering group or committee that sets principles and approves high-risk use cases.
- A lightweight risk-assessment checklist for every AI initiative.
- Clear ownership for each production AI system: technical owner and business owner.

8. A 90-day action plan to get started

You don't need a multi-year program to begin. Here is a realistic 90-day plan.

Weeks 0–2: Align and prioritize

- Identify 3–5 core business metrics you want to improve with AI and data.

- Run short workshops with frontline teams to collect pain points and ideas.
- Build and prioritize a use-case backlog using impact vs. feasibility.
- Select 1–2 flagship use cases for pilots.

Weeks 2–6: Design and prepare

- Define the data requirements and check availability and quality.
- Sketch a simple architecture for data flows and model integration.
- Agree on success metrics, guardrails, and governance requirements.
- Stand up or refine core pieces of your data and MLOps platform if needed.

Weeks 6–12: Pilot and learn

- Build an end-to-end pilot for the selected use cases.
- Deploy to a limited scope (e.g., one region, one team, or a subset of customers).
- Monitor business metrics, user feedback, and model performance.
- Decide whether to scale, iterate, or sunset based on evidence.

At the end of 90 days, you should have:

- At least one AI-powered workflow running in the real world.

- A clearer view of the data and platform gaps that matter most.
- A repeatable framework for your next set of AI initiatives.

9. How Salt Technologies can help

Salt Technologies is a software consulting and development partner focused on building practical, production-ready AI and data solutions.

We help organizations to:

- Turn scattered AI experiments into a clear, outcome-driven roadmap.
- Build or modernize data and MLOps foundations on cloud-native platforms.
- Design and deliver LLM-powered assistants, copilots, and analytics tools that your teams actually use.
- Implement responsible AI governance so you can innovate confidently.

Whether you're just starting with AI or looking to scale your existing efforts, we can support you across strategy, architecture, and implementation.

To explore what an AI & data strategy could look like for your organization, reach out to the Salt Technologies team.

https://www.salttechno.com/ | sales@salttechno.com